

Effect-Dependent Transformations for Concurrent Programs

Nick Benton

Microsoft Research, Cambridge, UK
nick@microsoft.com

Martin Hofmann

LMU, Munich, Germany
hofmann@ifi.lmu.de

Vivek Nigam

UFPB, João Pessoa, Brazil
vivek.nigam@gmail.com

Abstract

We describe a denotational semantics for an abstract effect system for a higher-order, shared-variable concurrent programming language. We prove the soundness of a number of general effect-based program equivalences, including a parallelization equation that specifies sufficient conditions for replacing sequential composition with parallel composition. Effect annotations are relative to abstract locations specified by contracts rather than physical footprints allowing us in particular to show the soundness of some transformations involving fine-grained concurrent data structures, such as Michael-Scott queues, that allow concurrent access to different parts of mutable data structures.

Our semantics is based on refining a trace-based semantics for first-order programs due to Brookes. By moving from concrete to abstract locations, and adding type refinements that capture the possible side-effects of both expressions and their concurrent environments, we are able to validate many equivalences that do not hold in an unrefined model. The meanings of types are expressed using a game-based logical relation over sets of traces. Two programs e_1 and e_2 are logically related if one is able to solve a two-player game: for any trace with result value v_1 in the semantics of e_1 (challenge) that the player presents, the opponent can present an (response) equivalent trace in the semantics of e_2 with a logically related result value v_2 .

1. Introduction

Type-and-effect systems refine conventional types with extra static information capturing a safe upper bound on the possible side-effects of expression evaluation. Since their introduction by Gifford and Lucassen [16], effect systems have been used for many purposes, including region-based memory management [11], tracking exceptions [21, 23], communication behaviour [5] and atomicity [15] for concurrent programs, and information flow [12].

A major reason for tracking effects is to justify program transformations, most obviously in optimizing compilation [9]. For example, one may remove computations whose results are unused, *provided* that they are sufficiently pure, or commute two state-manipulating computations, *provided* that the locations they may read and write are suitably disjoint. Several groups have recently studied the semantics of effect systems, with a focus on formally justifying such effect-dependent equational reasoning [6, 8, 10, 17, 25]. A common approach, which we follow here, is to interpret effect-refined types using a logical relation over the (denotational or operational) semantics of the unrefined (or untyped) language, simultaneously identifying both the subset of computations that have a particular effect type and a coarser notion of equivalence (or approximation) on that subset. Such a semantic approach decouples the meaning of effect-refined types from particular syntactic rules: one may establish that a term has a type using various more or less approximate inference systems, or by detailed semantic reasoning.

For sequential computations with global state, denotational models already provide significant abstraction. For example, the denotations of `skip` and `X++`; `X--` are typically equal, so it is immediate that the second is semantically pure. More generally, the meaning of a judgement $\Gamma \vdash e : \tau \& \varepsilon$ guarantees that the result

of evaluating e will be of type τ with side-effects at most ε , under assumptions Γ (a ‘rely’ condition), on the behaviour of e ’s free variables. The possible interaction points between e and its environment are restricted to initial states and parameter values, and final states and results, of e itself and its explicitly-listed free variables. Furthermore, all those interaction points are visible in the term and are governed by specific annotations appearing in the typing judgement.

For shared-variable concurrency, there are many more possible interactions. An expression’s environment now also includes anything that may be running concurrently and, moreover, atomic steps of e and its concurrent environment may be arbitrarily interleaved, so it is no longer sufficient to just consider initial and final states. A priori, this leads to far fewer equations between programs. For example, `X++`; `X--` may be distinguished from `skip` by being run concurrently with a command that reads or writes `X`. But few programs do anything useful in the presence of unconstrained interference, so we need ways to describe and control it. Fine-grained, optimistic algorithms, which rely on custom protocols being followed by multiple threads with concurrent access to a shared data structure, can significantly outperform ones based on coarse-grained locking, but are notoriously challenging to write and verify.

There is a huge literature on shared-variable concurrency, from type systems ensuring race-freedom of programs with locks [1] to sophisticated semantic models for reasoning about refinement of fine-grained concurrent datastructures [27]. This paper explores effect types as a straightforward, lightweight interface language for modular reasoning about equivalence and refinement, e.g. for safely transforming sequential composition into parallelism. We show how the semantics of a simple effect system scales smoothly to the concurrent setting, allowing us to control interference and prove non-trivial equivalences, extending (somewhat to our surprise) to the correctness of some fine-grained algorithms.

We build on a trace semantics for concurrent programs, due to Brookes [13], which explicitly describes possible interference by the environment. We extend Brookes’s semantics to a higher-order language and then refine it by a semantically-formulated effect system that separately tracks: (1) the store effects of an expression during evaluation; (2) the assumed effects of transitions by the environment; and (3) the overall end-to-end effect. Rather than tracking effects at the level of individual concrete heap cells, we view the heap as a set of abstract data structures, each of which may span several locations, or parts of locations [6]. Each abstract location has its own notion of equality, and its own notion of legal mutation. Write effects, for example, need only be flagged when the equivalence class of an abstract location may change. Both typing and refinement judgements may be established by a combination of generic type-based rules and semantic reasoning in the model.

We begin with some motivating examples.

Equivalence modulo non-interference: Our semantics justifies the following equation *at* the effect type `unit & {coX}` | ε | $\varepsilon \cup \{rd_X, wr_X\}$:

$$(X := !X + 1; X := !X + 1) = (X := !X + 2)$$

This says that the two commands are equivalent with return type unit , exhibit the effect co_X , signifying concurrent or ‘chaotic’ access to X along the way, and have an overall end-to-end effect of ε plus reading and writing X , *provided* that the effect, ε , of the concurrent environment does not involve X .

Overlapping References: Let p, p^{-1} implement a bijection $\mathbb{Z} \rightarrow \mathbb{Z} \times \mathbb{Z}$, and consider the following functions:

```
readFst () = p(!X).1
readSnd () = p(!X).2
wrtFst n = let rec try () = (let m = !X in let (x, y) = p(m) in
  let m' = p-1(n, y) in if cas(X, m, m') then () else try ())
  in try ()
wrtSnd n = let rec try () = let m = !X in let (x, y) = p(m) in
  let m' = p-1(x, n) in if cas(X, m, m') then () else try ()
  in try ()
```

which multiplex two abstract integer references onto a single concrete one. Note that the write functions, wrtFst and wrtSnd , use compare-and-swap, cas , to atomically update the value of the reference.

Our generic rules then say that a program, e_1 , that only reads and/or writes one abstract reference can be commuted, or executed in parallel, with another program, e_2 , that only reads and/or writes into a different reference. This lets one use types to, say, justify parallelizing a call to wrtFst followed by one to wrtSnd , even though they read and write the same concrete location, which looks like a race.

Version numbers: One can isolate a transaction that reads and then writes a piece of state simply by enclosing the whole thing in $\text{atomic}(\cdot)$. A more concurrent alternative adds a monotonic version number to the data. A transaction then works on a private copy, only committing its changes back (and incrementing the version) if the current version number is the same as that of the original copy. We can define an abstract integer reference \mathfrak{X} in terms of two concrete ones, X_{ver} and X_{val} , governed by a specification that says $!X_{\text{val}}$ may only change when $!X_{\text{ver}}$ increases. We define

```
transact f = let rec try() = let (val, ver) = atomic(!Xval, !Xver)
  in let res = f(val) in if atomic(if !Xver = ver then
    Xver := ver + 1; Xval := res; true else false)
  then () else try()
  in try()
```

Under the assumption that f is a pure function (has effect type $\text{int} \xrightarrow[\varepsilon]{0|\varepsilon} \text{int}$ for any ε), we can show

$$\text{transact } f = \text{atomic}(X_{\text{val}} := f(!X_{\text{val}}); X_{\text{ver}} := !X_{\text{ver}} + 1)$$

at type $\text{unit} \& \{rd_{\mathfrak{X}}, wr_{\mathfrak{X}}\} \mid \varepsilon \mid \varepsilon \cup \{rd_{\mathfrak{X}}, wr_{\mathfrak{X}}\}$ for any ε not including chaotic access, $\text{co}_{\mathfrak{X}}$, to \mathfrak{X} . The environment effect ε here *may* include reading and writing \mathfrak{X} , so concurrent calls to transact are linearizable.

Loop Parallelization: Our next example is inspired by a loop unrolling optimization [26]. Assume given a linked list of integers




Figure 1. Illustration of a Michael-Scott Queue. The list resulting from the pointer to the element n_0 (the head pointer with the continuous arrow in black) contains the list of elements $[n_1, \dots, n_j]$. The enqueueing operation is illustrated by the dotted arrow and the box with the element n_{j+1} (in blue), while the dequeueing operation is illustrated by the dot dashed head pointer (in red).

```
dequeue () = let rec try () =
  let n0 = !head in if !n0.next = null then null
  else let n1 = !n0.next in
    if cas(!head, n0, n1) then !n1.ele else try ()
  in try ()
enqueue(x) = let rec try (p) =
  if !p.next = null then
    if atomic(if !p.next = null then
      !p.next := ref(x, null); true else false)
    then () else try (!p.next)
  else try (!p.next)
  in try (!head)
```

Figure 2. Enqueue and Dequeue programs for a Michael-Scott Queue at location head .

pointed by head . Consider the following functions:

```
map f = let rec applyf n =
  n.ele := f(n.ele); if n.next = null then ()
  else applyf (n.next)
  in if !head = null then () else applyf (!head)
map2Par f = let rec applyf2 n =
  n.ele := f(n.ele) || n.next.ele := f(n.next.ele);
  if n.next.next = null then ()
  else if n.next.next.next = null then
    n.next.next.ele := f(n.next.next.ele)
    else applyf2 (n.next.next)
  in if !head = null then ()
  else if !head.next = null then
    !head.next.ele := f(!head.next.ele)
  else applyf2 (!head)
```

The function map simply applies a pure function f to each element of the list, each element per iteration. The function map2Par , on the other hand, applies f to two consecutive elements of the list in parallel, potentially allowing one to exploit multiple cores. Our effect-based reasoning will soundly transform map into map2Par (under the assumption that the environment does not interfere with the list).

Michael-Scott Queue: The Michael-Scott Queue [20] (MSQ) is a fine grained concurrent data structure, allowing threads to access and modify different parts of a queue safely and simultaneously. We present a version like that of Turon et al [27], which is an idealized version of the MSQ, without a tail pointer.

An MSQ maintains a pointer *head* to a non-empty linked list as depicted in Figure 1. The first node, the node containing the element n_0 in the figure, is not an element of the queue, but is a “sentinel”. Hence the queue in the figure holds $[n_1, \dots, n_j]$.

The enqueue and dequeue operations are defined in Figure 2 and illustrated in Figure 1. Elements are dequeued from the beginning of the linked list, and enqueued at the end, which involves a traversal that is done without locking. Once the end, p , of the linked list is found, the program atomically attempts to insert the new element. This is necessary because other programs may have enqueued elements to the end of the list, meaning that p is no longer the end of the list.

The dequeue operation should move the *head* pointer from the current sentinel, n_0 , to the following element n_1 . However, as other programs may also be attempting to dequeue an element, we use compare-and-swap to atomically update the *head* pointer if *head* still points to the same sentinel. Notice that the dequeued elements can still reach the sentinel of the queue. (In Figure 1, these are the nodes containing n_{-k}, \dots, n_{-1} .) This is necessary because there might be other (slower) threads that want to enqueue an element and are still searching for the end of the list by traversing the portion of the queue that has already been dequeued.

We prove that the enqueue and dequeue of Figure 2 are equivalent to `atomic(enqueue)` and `atomic(dequeue)`, their atomic versions which perform all operations in a single step, at a type that allows the environment to be concurrently reading and writing the queue. So the fine-grained MSQ behaves like a synchronized queue, as might also be implemented using locks.

2. Syntax

In this section we define the syntax of a metalanguage for concurrent, stateful computations and higher-order functions. Communication between parallel computations is via a shared heap mapping dynamically allocated locations to structured values, which include pointers. To keep the model simple, we do not allow functions to be stored in the heap (no higher-order store).

Memory model We assume a countably infinite set \mathbb{L} of physical locations X_1, \dots, X_n, \dots and a set \mathbb{VB} of “R-values” that can be stored in those references including integers, booleans, locations, and tuples of R-values, written (v_1, \dots, v_n) . We assume that it is possible to tell of which form a value is and to retrieve its components in case it is a tuple. A heap h , then, is a *finite map* from \mathbb{L} to \mathbb{VB} , written $\{(X_1, c_1), (X_2, c_2), \dots, (X_n, c_n)\}$, specifying that the value stored in location X_i is c_i . We write $\text{dom}(h)$ for the domain of h and write $h[X \mapsto c]$ for the heap that agrees with h except that it gives the variable X the value c . The set of heaps is denoted by \mathbb{H} . We also assume that $\text{new}(h, v)$ yields a pair (X, h') where $X \in \mathbb{L}$ is a fresh location and $h' \in \mathbb{H}$ is $h[X \mapsto v]$.

Syntax of expressions The syntax of untyped values and computations is:

$$\begin{aligned} v &::= x \mid (v_1, v_2) \mid v_r \mid c \mid \text{rec } f \ x = t \\ e &::= v \mid \text{let } x = e_1 \text{ in } e_2 \mid v_1 \ v_2 \mid \text{if } v \text{ then } e_1 \text{ else } e_2 \\ &\quad \mid !v \mid v_1 := v_2 \mid \text{ref}(v) \mid e_1 \parallel e_2 \mid \text{atomic}(e) \end{aligned}$$

Here, x ranges over variables, v_r over R-values, and c over built-in functions, which include arithmetic, testing whether a value is an integer, function, pair or reference, equality on simple values, etc. Each c has a corresponding semantic partial function F_c , so for example $F_+(n, n') = n + n'$ for integers n, n' .

The construct `rec f $x = e$` defines a recursive function with body e and recursive calls made via f ; we use $\lambda x.e$ as syntactic sugar in the case when f is not free in e . Next, `! v` (reading) returns the contents of location v , $v_1 := v_2$ (writing) updates location v_1 with value v_2 , and `ref(v)` (allocating) returns a fresh location initialized with v . The metatheory is simplified by using “let-normal form”,

in which the only elimination for computations is `let`, though we sometimes nest computations as shorthand for let-expanded versions in examples.

The construct $e_1 \parallel e_2$ is evaluated by arbitrarily interleaving evaluation steps of e_1 and e_2 until each has produced a value, say v_1 and v_2 ; the result is then (v_1, v_2) . Assignment, dereferencing and allocation are atomic, but evaluation of nested expressions is generally not. To enforce atomicity, `atomic(e)` evaluates an arbitrary e in one step, without any environmental interference. One can then define a (more realistic) compare-and-swap operation `cas(X, v_1, v_2)`:

$$\text{cas}(X, v_1, v_2) = \text{atomic}(\text{if } !X = v_1 \text{ then } X := v_2; \text{true} \text{ else } \text{false})$$

this atomically both checks if location X contains v_1 and, if so, replaces it with v_2 and returns `true`; otherwise the location is unchanged and the returned value is `false`.

We define the free variables, $FV(e)$, of a term, closed terms, and the substitution $e[v/x]$ of v for x in e , in the usual way. Locations may occur in terms, but the type system will constrain their use.

3. Denotational Model

We now sketch a denotational semantics for our metalanguage based on Brookes’ trace semantics [13]. Fuller details can be found in a technical report (attached), which in particular establishes computational adequacy of the model with respect to a small-step operational semantics using interleaving.

3.1 Preliminaries

A *predomain* is an ω -cpo, i.e., a partial order with suprema of ascending chains. A *domain* is a predomain with a least element, \perp . Recall that $f : A \rightarrow A'$ is *continuous* if it is monotone $x \leq y \Rightarrow f(x) \leq f(y)$ and preserves suprema of chains, i.e., $f(\sup_i x_i) = \sup_i f(x_i)$. Any set is a predomain with the discrete order (flat predomain). If X is a set and A a predomain then any $f : X \rightarrow A$ is continuous. We denote a partial (continuous) function from set (predomain) A to set (predomain) B by $f : A \rightarrow B$. If A, B are predomains the cartesian product $A \times B$ and the set of continuous functions $A \rightarrow B$ form themselves predomains (with the obvious componentwise and pointwise orders) and make the category of predomains cartesian closed. Likewise, the partial continuous functions $A \rightarrow B$ between predomains A, B form a domain.

If $P \subseteq A$ and $Q \subseteq B$ are subsets of predomains A and B we define $P \times Q \subseteq A \times B$ and $P \rightarrow Q \subseteq A \rightarrow B$ in the usual way. We may write $f : P \rightarrow Q$ for $f \in P \rightarrow Q$.

A subset $U \subseteq A$ is *admissible* if whenever $(a_i)_i$ is an ascending chain in A such that $a_i \in U$ for all i , then $\sup_i a_i \in U$, too. If $f : X \times A \rightarrow A$ is continuous and A is a domain then one defines $f^\sharp(x) = \sup_i f_i^\sharp(\perp)$ with $f_i(a) = f(x, a)$. One has, $f(x, f^\sharp(x)) = f^\sharp(x)$ and if $U \subseteq A$ is admissible and contains \perp and $f : X \times U \rightarrow U$ then $f^\sharp : X \rightarrow U$, too. An element d of a predomain A is *compact* if whenever $d \leq \sup_i a_i$ then $d \leq a_i$ for some i . E.g. in the domain of partial functions from \mathbb{N} to \mathbb{N} the compact elements are precisely the finite ones. A continuous partial function $f : A \rightarrow A$ is a *retract* if $f(a) \leq a$ and $f(f(a)) = f(a)$ hold for all $a \in A$. In short: $f \leq \text{id}_A$ and $f \circ f \leq f$. If, in addition, f has a finite image then f is called a *deflation* [3]. Note that if f is a retract then $\text{dom}(f) = \text{Img}(f)$ and if $a \in \text{Img}(f)$ then $a = f(a)$. We also note that if a is in the image of a deflation then a is compact.

We define the usual state monad on predomains, by taking $SA = \mathbb{H} \rightarrow \mathbb{H} \times A$.

Definition 3.1. Let P be a subset of a predomain A . Then $\text{Adm}(P)$ is the least admissible superset of P . Concretely, $a \in \text{Adm}(P)$ iff there exists a chain $(a_i)_i$ such that $a_i \in P$ for all i and $a = \sup_i a_i$.

Lemma 3.2. If $f : A_1 \times \dots \times A_n$ is continuous; $P_i \subseteq A_i$ are arbitrary subsets and $Q \subseteq B$ is admissible then $f : P_1 \times \dots \times P_n \rightarrow Q$ implies $f : \text{Adm}(P_1) \times \dots \times \text{Adm}(P_n) \rightarrow Q$.

Lemma 3.3. *Let A, B be predomains and let $(p_i)_i$ be a chain of retracts on B such that $p_i(b)$ is compact for each i and $\sup_i p_i = \text{id}_B$ and $b \in Q$ implies $p_i(b) \in Q$ for all i . Then $P \rightarrow \text{Adm}(Q) = \text{Adm}(P \rightarrow Q)$.*

3.2 Traces

A trace models a terminating run of a concurrent computation as a sequence of pairs of heaps, each representing pre- and post-state of one or more atomic actions. The semantics of a program then is a (typically large) set of traces (and final values), accounting for all possible environment interactions.

Definition 3.4 (Traces). *A trace is a finite sequence of the form $(h_1, k_1)(h_2, k_2) \dots (h_n, k_n)$ where for $1 \leq j \leq i \leq n$, we have $h_i, k_i \in \mathbb{H}$ and $\text{dom}(h_j) \subseteq \text{dom}(h_i)$, $\text{dom}(h_j) \subseteq \text{dom}(k_i)$, $\text{dom}(k_j) \subseteq \text{dom}(h_i)$, $\text{dom}(k_j) \subseteq \text{dom}(k_i)$. We write Tr for the set of traces.*

Let t be a trace. A trace of the form $u(h, h)v$ where $t = uv$ is said to arise from t by stuttering. A trace of the form $u(h, k)v$ where $t = u(h, q)(q, k)v$ is said to arise from t by mumbling. For example, if $t = (h_1, k_1)(h_2, k_2)(h_3, k_3)$ then $(h_1, k_1)(h, h)(h_2, k_2)(h_3, k_3)$ arises from t by stuttering. In the case where $k_1 = h_2$ the trace $(h_1, k_2)(h_3, k_3)$ arises from t by mumbling. A set of traces U is closed under stuttering and mumbling if whenever t' arises from t by stuttering or mumbling and $t \in U$ then $t' \in U$, too.

Brookes [13] gives a fully-abstract semantics for while-programs with parallel composition using sets of traces closed under stuttering and mumbling. We here extend his semantics to higher-order functions and general recursion.

Definition 3.5 (Trace Monad). *Let A be a predomain. Elements of the domain TA are sets U of pairs (t, a) where t is a trace and $a \in A$ such that the following properties are satisfied:*

- $[S\&M]$: if t' arises from t by stuttering or mumbling and $(t, a) \in U$ then $(t', a) \in U$.
- $[Down]$: if $(t, a_1) \in U$ and $a_2 \leq a_1$ then $(t, a_2) \in U$.
- $[Sup]$: if $(a_i)_i$ is a chain in A and $(t, a_i) \in U$ for all i then $(t, \sup_i a_i) \in U$.

The elements of TA are partially ordered by inclusion.

Lemma 3.6. *If A is a predomain then TA is a domain.*

An element U of TA represents the possible outcomes of a nondeterministic, interactive computation with final result in A . Thus, if $(t, a) \in U$ for $t = (h_1, k_1) \dots (h_n, k_n)$ then there could be n interactions with the environment with heaps h_1, \dots, h_n being “played” by the environment and “answered” with heaps k_1, \dots, k_n by the computation. After that, this particular computation ends and a is the final result value.

For example, the semantics of $X := !X + 1; X := !X + 1; !X$ contains many traces, including the following, where we write $[n]$ for the heap in which X has value n :

$(([10], [12]), 12),$
 $(([10], [11])([15], [16]), 16),$
 $(([10], [11])([15], [16])([17], [17]), 17),$
 $(([10], [11])([15], [16])([17], [17]), 16),$
 $(([10], [11])([17], [17])([15], [16]), 16), \dots$

Axiom $[S\&M]$ is taken from Brookes. It ensures that the semantics does not distinguish between late and early choice [27] and related phenomena which are reflected, e.g., in resumption semantics [24], but do not affect observational equivalence. Note that non-termination is modelled by the empty set, so we are working with an ‘angelic’ notion of equivalence (‘may semantics’ [22]). For example, the semantics of $X := 0; \text{if } X=0 \text{ then } 0 \text{ else diverge}$ is the same as that of $X := 0; 0$ and contains, for example $(([10], [0]), 0)$ but also (stuttering) $(([10], [0]), ([34], [34]), 0)$.

Note that it is not possible to tell from a trace whether an external update of X has happened before or after the reading of X .

Let us also illustrate how traces iron out some intensional differences that show up when concurrency is modelled using transition systems or resumptions. Consider the following two programs where $?$ denotes a nondeterministically chosen boolean value.

$e_1 \equiv \text{if } ? \text{ then } X := 0; \text{true else } X := 0; \text{false}$
 $e_2 \equiv X := 0; ?$

Both e_1 and e_2 admit the same traces, namely $(([x], [0]), \text{true})$ and $(([x], [0]), \text{false})$ and stuttering variants thereof. In semantic models based on transition systems or resumptions and bisimulation, these are distinguished, which necessitates the use of special mechanisms such as history and prophecy variables [2], forward-backward simulation [19], or speculation [27] in reasoning.

Axioms $[Down]$ and $[Sup]$ are known from the Hoare powerdomain [24]. Recall that the Hoare powerdomain PA contains the subsets of A which are downclosed ($[Down]$) and closed under suprema of chains ($[Sup]$). Such subsets are also known as Scott-closed sets. Thus, TA is the restriction of $P(\text{Tr} \times A)$ to the sets closed under stuttering and mumbling. Axiom $[Down]$ ensures that the ordering is indeed a partial order and not merely a preorder. Additional nondeterministic outcomes that are less defined than existing ones are not recorded in the semantics.

Definition 3.7. *If $U \subseteq \text{Tr} \times A$ then U^\dagger is the least subset of TA containing U , i.e. U^\dagger is the closure of U under $[S\&M]$, $[Down]$, $[Sup]$.*

Definition 3.8. *Let A, B be a predomains. We define the continuous functions $\text{rtn} : A \rightarrow TA$ and $\text{bnd} : (A \rightarrow TB) \times TA \rightarrow TB$ by:*

$$\text{rtn}(a) := \{((h, h), a) \mid h \in \mathbb{H}\}^\dagger$$

$$\text{bnd}(f, g) := \{(\{uv, b\} \mid (u, a) \in g \wedge (v, b) \in f(a))\}^\dagger$$

These endow TA with the structure of a strong monad. The continuous function $\text{fromstate} : SA \rightarrow TA$ is defined by:

$$\text{fromstate}(c) := \{((h, k), a) \mid c(h) = (k, a)\}^\dagger$$

If t_1, t_2, t_3 are traces, we write $\text{inter}(t_1, t_2, t_3)$ to mean that t_3 can be obtained by interleaving t_1 and t_2 in some way, i.e., t_3 is contained in the shuffle of t_1 and t_2 . In order to model parallel composition we introduce the following helper function

$$\mid : TA \times TB \rightarrow T(A \times B)$$

$$U \mid V := \{(t_3, (a, b)) \mid \text{inter}(t_1, t_2, t_3), (t_1, a) \in U, (t_2, b) \in V\}^\dagger$$

The continuous map $\text{at} : TA \rightarrow TA$ is defined by:

$$\text{at}(U) := \{((h, k), v) \mid ((h, k), v) \in U\}^\dagger$$

Notice that due to mumbling $((h, k), v) \in U$ iff there exists an element $((h_1, h_2)(h_2, h_3) \dots (h_{n-2}, h_{n-1})(h_{n-1}, h_n), v) \in U$ where $h = h_1$ and $h_n = k$. The presence of such an element, however, models an atomic execution of the computation represented by U .

3.3 Semantic values

The predomain \mathbb{V} of untyped values is the least solution of the following domain equation:

$$\mathbb{V} \simeq \mathbb{V}\mathbb{B} + (\mathbb{V} \rightarrow T\mathbb{V}) + \mathbb{V}^*.$$

That is, values are either R-values, continuous functions from values to computations ($T\mathbb{V}$), or tuples of values. We tend to identify the summands of the right hand side with subsets of \mathbb{V} but may use tags like $\text{fun}(f) \in \mathbb{V}$ when $f : \mathbb{V} \rightarrow T\mathbb{V}$ to avoid ambiguity.

We have families of deflations $p_i : \mathbb{V} \rightarrow \mathbb{V}$ and $q_i : T\mathbb{V} \rightarrow T\mathbb{V}$, referred to as canonical deflations, so that $(p_i)_i$ and $(q_i)_i$ are ascending chains converging to the identity. The definition is entirely standard and may be found in the accompanying material. It shows in particular that \mathbb{V} and $T\mathbb{V}$ are *bifinite* (equivalently SFP) (pre-)domains [3] and as such also Scott (pre-) domains. The presence

of these deflations allows us to apply Lemma 3.3 and simplifies reasoning in general.

The semantics of values $\llbracket v \rrbracket \in \mathbb{V} \rightarrow \mathbb{V}$ and terms $\llbracket t \rrbracket \in \mathbb{V} \rightarrow T\mathbb{V}$ are given by the recursive clauses in Figure 3. Environments, ρ , are properly tuples of values; we abuse notation slightly by treating them as maps from variables, x , to values, v , (and write $\rho[x \mapsto v]$ for functional update) to avoid mentioning an explicit context in which untyped terms are well-formed.

4. Abstract Locations

We build on the concept of abstract locations defined by Benton, Hofmann, and Nigam [6]. These allow complicated data structures that span several concrete locations, or only parts of them, to be regarded as a single “location” that can be written to and read from. Essentially, an abstract location is given by a partial equivalence relation on heaps modelling well-formedness and equality together with a transitive relation modelling allowed modifications of the abstract location. Abstract locations then allow certain commands that modify the physical heap to be treated as read-only or even pure if they respect the contracts. Abstract locations are related to *islands* [4] which also allow one to specify heap allocated data structures and use transition systems for that purpose. An important difference is that abstract locations do not require physical footprints in the form of sets of concrete locations.

Due to the absence of dynamic allocation at the level of abstract locations in the present paper, we can slightly simplify the original definition [6], dropping those axioms that involve the interaction with dynamic allocation.¹ On the other hand, in the presence of concurrency, we need *two* partial equivalence relations: one that models semantic equivalence and well-formedness and a finer one that constrains the heap modifications that other concurrent computations that are independent of the given abstract locations are allowed to do *while* an operation on the abstract location is ongoing, but temporarily preempted.

Definition 4.1 (Concurrent Abstract Location). A concurrent abstract location \mathbb{L} consists of the following data:

(1) a partial equivalence relation \sim on \mathbb{H} modeling the “semantic equivalence” on the bits of the store that \mathbb{L} uses. If $h \sim h'$ then the same computation started on h and h' , respectively, will yield related or even equal results.

(2) a partial equivalence relation \equiv on \mathbb{H} refining \sim and modeling the “strict equivalence” on the bits of the store that \mathbb{L} uses. If a concurrent computation on \mathbb{L} has reached h and is preempted, then another computation may replace h with h' where $h \equiv h'$ and then the original computation on \mathbb{L} may resume on h' without the final result being compromised.

(3) a transitive (and reflexive on the support of \sim) relation $\xrightarrow{\mathbb{L}}$ modeling how exactly the heap may change upon writing the abstract location and in particular what bits of the store such writes leave intact. In other words, if $h \xrightarrow{\mathbb{L}} h_1$ then h_1 might arise by writing to \mathbb{L} in h and all possible writes are specified by $\xrightarrow{\mathbb{L}}$. We call $\xrightarrow{\mathbb{L}}$ the step relation of \mathbb{L} .

In addition, we require the following conditions where $h : \mathbb{L}$ stands for $h \sim h$.

1. If $h : \mathbb{L}$ then $h \equiv h$;
2. if $h \xrightarrow{\mathbb{L}} h_1$ then $h : \mathbb{L}$ and $h_1 : \mathbb{L}$.

¹ Though our examples do all satisfy these axioms, leaving the way open to a future extension with dynamically allocation of abstract locations and concurrency.

If $h \xrightarrow{\mathbb{L}} h_1$ and at the same time $h \equiv h_1$, then we say that h_1 arises from h by a silent move in \mathbb{L} . Our semantic framework will permit silent moves at all times.

We now introduce some examples of abstract locations.

Single Integer For our simplest example, consider the following abstract location parametric with respect to concrete location X as follows:

$$\begin{aligned} h &\stackrel{\text{int}(X)}{\sim} h' && \iff \exists n. h(X) = \text{int}(n) \wedge h'(X) = \text{int}(n) \\ h &\stackrel{\text{int}(X)}{\equiv} h' && \iff h \stackrel{\text{int}(X)}{\sim} h' \\ h &\xrightarrow{\text{int}(X)} h_1 && \iff h : \text{int}(X), h_1 : \text{int}(X) \text{ and } \forall X' \in \mathbb{L}. X' \neq X \Rightarrow h(X') = h_1(X') \end{aligned}$$

Two heaps are semantically equivalent (w.r.t. $\text{int}(X)$ that is) if the values stored in X are integers and equal; the step relation requires all other concrete locations to be unchanged.

We will sometimes abuse notation and write rd_X, wr_X, co_X for $rd_{\text{int}(X)}, wr_{\text{int}(X)}, co_{\text{int}(X)}$.

Overlapping references Let X be a concrete location encoding a pair of integer values using a bijection p . We define the abstract location $\text{fst}(X)$ as below. We omit $\text{snd}(X)$ which is similar, but only looks at the second projection, instead of the first.

$$\begin{aligned} h &\stackrel{\text{fst}(X)}{\sim} h' && \iff \exists a_1 a_2 a'_1 a'_2 \in \mathbb{Z}. h(X) = p^{-1}(a_1, a_2) \wedge h'(X) = p^{-1}(a'_1, a'_2) \wedge a_1 = a'_1 \\ h &\stackrel{\text{fst}(X)}{\equiv} h' && \iff h \stackrel{\text{fst}(X)}{\sim} h' \\ h &\xrightarrow{\text{fst}(X)} h_1 && \iff h : \text{fst}(X), h_1 : \text{fst}(X) \text{ and } (\forall X' \neq X. h(X') = h_1(X')) \wedge (\forall a_1 a_2 a'_1 a'_2 \in \mathbb{Z}. h(X) = p^{-1}(a_1, a_2) \wedge h_1(X) = p^{-1}(a'_1, a'_2) \Rightarrow a_2 = a'_2) \end{aligned}$$

The semantic (and strict) equivalence of $\text{fst}(X)$ (respectively, $\text{snd}(X)$) specifies that two heaps h and h' are equivalent whenever they both store a pair of values in X and the first projections (respectively, second projection) of these pairs are the same. The step relation of $\text{fst}(X)$ (respectively, $\text{snd}(X)$) specifies that it keeps all other locations alone and does not change the second projection (respectively, first projection) of the pair stored at location X .

Version Numbers The abstract location \mathfrak{X} consists of two concrete locations X_{val} and X_{ver} and its relations are specified as follows:

$$\begin{aligned} h &\stackrel{\mathfrak{X}}{\sim} h' && \iff h(X_{\text{val}}) = h'(X_{\text{val}}) \\ h &\stackrel{\mathfrak{X}}{\equiv} h' && \iff h \stackrel{\mathfrak{X}}{\sim} h' \\ h &\xrightarrow{\mathfrak{X}} h_1 && \iff \forall X' \notin \{X_{\text{ver}}, X_{\text{val}}\}. h(X') = h_1(X') \wedge h : \mathfrak{X} \wedge h_1 : \mathfrak{X} \wedge h(X_{\text{ver}}) \leq h_1(X_{\text{ver}}) \wedge [h(X_{\text{val}}) \neq h_1(X_{\text{val}}) \Rightarrow h(X_{\text{ver}}) < h_1(X_{\text{ver}})] \end{aligned}$$

Two heaps are semantically equivalent if they have the same value (independent of the version number). The step relation specifies that the version number does not decrease and it increases if the value changes.

Loop Parallelization For a concrete location X , we introduce two concurrent abstract locations $\text{listeven}(X)$ and $\text{listodd}(X)$, which only look, respectively, at the elements in the even and odd positions of the linked list pointed by X . Formally, let $L(X, h)$ denote that $h(X)$ points to a well formed linked list of integers of length $L(X, h).len$ and locations $L(X, h).locs$ and that $L(X, h)[i]$ is the i^{th} node of the list for $1 \leq i \leq L(X, h).len$. The relations for $\text{listeven}(X)$ are as below. We omit the relations for $\text{listodd}(X)$, which

$\llbracket x \rrbracket \rho = \rho(x)$	$\llbracket v \rrbracket \rho = \text{rtn}(\llbracket v \rrbracket \rho)$
$\llbracket v_r \rrbracket \rho = v_r$	$\llbracket \text{let } x = e_1 \text{ in } e_2 \rrbracket \rho = \text{bnd}(\lambda d. \llbracket e_2 \rrbracket \rho[x \mapsto d], \llbracket e_1 \rrbracket \rho)$
$\llbracket (v_1, v_2) \rrbracket \rho = (\llbracket v_1 \rrbracket \rho, \llbracket v_2 \rrbracket \rho)$	$\llbracket v_1 \ v_2 \rrbracket \rho = \llbracket v_1 \rrbracket \rho(\llbracket v_2 \rrbracket \rho)$
$\llbracket v.i \rrbracket \rho = d_i \text{ if } i = 1, 2, \llbracket v \rrbracket \rho = (d_1, d_2)$	$\llbracket \text{if } v \text{ then } e_1 \text{ else } e_2 \rrbracket \rho = \llbracket e_1 \rrbracket \rho, \text{ if } \llbracket v \rrbracket \rho = \text{true}$
$\llbracket c \rrbracket \rho = \text{fun}(f)$	$\llbracket \text{if } v \text{ then } e_1 \text{ else } e_2 \rrbracket \rho = \llbracket e_2 \rrbracket \rho, \text{ if } \llbracket v \rrbracket \rho = \text{false}$
where $f(v) = \text{rtn}(F_c(v))$ if $F_c(v)$ is defined	$\llbracket !v \rrbracket \rho = \text{fromstate}(\lambda h. (h, h(X))), \text{ when } \llbracket v \rrbracket \rho = X$
and $f(v) = \emptyset$, otherwise.	$\llbracket v_1 := v_2 \rrbracket \rho = \text{fromstate}(\lambda h. (h[X \mapsto \llbracket v_2 \rrbracket \rho], \emptyset)), \text{ if } \llbracket v_1 \rrbracket \rho = X$
$\llbracket \text{rec } f \ x = e \rrbracket \rho = \text{fun}(g^{\ddagger}(\rho))$	$\llbracket \text{ref}(v) \rrbracket \rho = \text{fromstate}(\lambda h. \text{new}(h, \llbracket v \rrbracket \rho))$
where $g(\rho, u) = \lambda d. \llbracket e \rrbracket \rho[f \mapsto u, x \mapsto d]$	$\llbracket \text{atomic}(e) \rrbracket \rho = \text{at}(\llbracket e \rrbracket \rho)$
$\llbracket v \rrbracket \rho = 0$, otherwise	$\llbracket e_1 \rrbracket \rho \mid \llbracket e_2 \rrbracket \rho = \llbracket e_1 \rrbracket \rho \mid \llbracket e_2 \rrbracket \rho$
	$\llbracket e \rrbracket \rho = \emptyset$, otherwise

Figure 3. Denotational semantics

are similar.

$$\begin{aligned}
h \stackrel{\text{listen}(X)}{\sim} h' &\iff L(X, h) \wedge L(X, h') \wedge L(X, h).len = L(X, h').len \wedge \\
&\quad L(X, h)[2i] = L(X, h')[2i] \\
&\quad \text{for } 0 \leq i \leq \lfloor L(X, h).len/2 \rfloor \\
h \stackrel{\text{listen}(X)}{=} h' &\iff h \stackrel{\text{listen}(X)}{\sim} h' \\
h \xrightarrow{\text{listen}(X)} h_1 &\iff h : \text{listen}(X) \wedge h_1 : \text{listen}(X) \wedge \\
&\quad L(X, h) \wedge L(X, h_1) \wedge \text{for } 0 \leq i \leq \lfloor L(X, h).len/2 \rfloor \\
&\quad L(X, h)[2i+1] = L(X, h_1)[2i+1] \wedge \\
&\quad L(X, h)[2i].next = L(X, h_1)[2i].next \wedge \\
&\quad \forall X' \notin L(X, h).locs. h(X') = h_1(X')
\end{aligned}$$

Michael-Scott queue For concrete location X we introduce a concurrent abstract location $\text{msq}(X)$ first informally as follows: we have $h \stackrel{\text{msq}(X)}{\sim} h'$ if both h and h' contain a well-formed MSQ rooted at X and these queues contain the same entries in the same order. They may, however, use different locations for the nodes and also have different garbage tails.

The relation $h \stackrel{\text{msq}(X)}{=} h'$ asserts that h and h' are identical on the part reachable and co-reachable from X via *next* pointers. This means that while an MSQ operation is working on the queue no concurrent operation working elsewhere is allowed to relocate the queue or remove the garbage trail which would be the case if we merely required that such operations do not change the $\text{MSQ}(X)$ -class.

The relation $h \xrightarrow{\text{msq}(X)}$, finally, is defined as the transitive closure of the actions of operations on the MSQ: adding nodes at the tail and moving nodes from the head to the garbage tail.

We now give a formal definition. We represent pointers *head*, *next*, *elem* using some layout convention, e.g. $v.\text{head} = v.1$, etc. We then define

$$h, X \xrightarrow{\text{next}} X' \iff X' \text{ can be reached from } X \text{ in } h \text{ by following a chain of next pointers}$$

We use $\text{List}(X, h, (X_0, \dots, X_n), (v_1, \dots, v_n))$ to signal that $h(X)$ points to a linked list with nodes X_0, \dots, X_n and entries v_1, \dots, v_n . Note that the first node X_0 acts as a sentinel and its *elem* component is ignored. Formally:

$$\begin{aligned}
h(X).head &= X_0 & h(X_i).elem &= v_i \text{ for } i = 1, \dots, n \\
h(X_i).next &= X_{i+1} \text{ for } i = 0, \dots, n-1 & h(X_n).next &= \text{null}
\end{aligned}$$

We define $\text{fp}(X, h)$ as the set of locations reachable and co-reachable from X via *next*, formally:

$$\text{fp}(X, h) = \{X' \mid X \xrightarrow{\text{next}} X' \vee X' \xrightarrow{\text{next}} X\}$$

Finally, we define $\text{snoc}(h, h', X, v)$ to mean that h' arises from h by attaching a new node containing v at the end of the list pointed to by X in h . Thus, in particular, $\text{List}(X, h, (X_0, \dots, X_n), (v_1, \dots, v_n))$ implies $\text{List}(X, h', (X_0, \dots, X_n, X_{n+1}), (v_1, \dots, v_n, v))$ for some $X_{n+1} \notin$

$\text{dom}(h)$. We omit the obvious frame conditions. We now define

$$\begin{aligned}
h \stackrel{\text{msq}(X)}{\sim} h' &\iff \exists \vec{X} \ \exists \vec{v}. \text{List}(X, h, \vec{X}, \vec{v}) \wedge \text{List}(X, h', \vec{X}', \vec{v}) \\
h \stackrel{\text{msq}(X)}{=} h' &\iff h \stackrel{\text{msq}(X)}{\sim} h' \wedge \forall X' \in \text{fp}(X, h). h(X') = h'(X') \\
h \xrightarrow{\text{msq}(X)} h_1 &\iff h : \text{msq}(X) \wedge h_1 : \text{msq}(X) \wedge \text{step}^*(h, h_1) \\
\text{step}(h, h_1) &\iff \forall X' \neq X. h(X') = h_1(X') \wedge \\
&\quad [h_1(X) = h(X).next \vee \exists v. \text{snoc}(h, h_1, X, v)]
\end{aligned}$$

In all of these examples, the only silent moves are identity moves. This is not so in the examples from [6] which contained data-structures that would reorganize during lookups and also patterns like late initialisation.

4.1 Worlds

We will group the abstract locations used to describe a program into a *world*. In this paper we do not model dynamic evolution of worlds; all abstract locations ever used must be set up upfront. While allocation of concrete locations may happen to increase a data structure modelled by an abstract location, e.g. in the Michael-Scott Queue example, no new such datastructures can appear. It is possible, however, to extend our work in this direction by using (proof-relevant) Kripke logical relations [4, 6].

Definition 4.2 (world). A world is a set of abstract locations.

The relation $h \models w$ (heap h satisfies world w) is defined as the largest relation such that $h \models w$ implies

- $h : l$ for all $l \in w$;
- if $l \in w$ and $h \xrightarrow{l} h_1$ then $h \stackrel{l}{=} h_1$ holds for all $l' \in w$ with $l' \neq l$ and $h_1 \models w$.

The original account of abstract locations [6] also has a notion of independence of locations which facilitates reasoning in the presence of dynamic allocation, and in particular permitted relocation of abstract locations. Since we are not currently treating dynamic allocation of abstract locations, we can avoid this notion here.

We remark that if our world w contains two obviously “dependent” abstract locations, e.g. has both an integer location and a boolean location placed at the same physical location, then there will be no heap h such that $h \models w$.

We assume a fixed *current* world w which may appear in definitions without being notationally reflected. See also Assumption 1.

5. Effects

For each abstract location l we have three elementary effects rd_l (reading from l), wr_l (writing to l), and co_l (chaotic or concurrent access). The chaotic access is similar to writing, but allows writes that are not in sync. For example, $e_1 = X := 1$ and $e_2 = X := 2$ both have individually the wr_X effect, but e_1 and e_2 are distinguishable with a context that assumes the wr_X -effect. Thus, e_1 and e_2 are not equal “at type” wr_X . At type co_X they are, however, equal, because a context that copes with this effect may not assume that both produce equal results.

We use the co_l effect to tell the environment not to look at a particular location during a concurrent computation. For example, we will be able to show that $X := !X + 1; X := !X + 1$ is equivalent to $X := !X + 2$ “at type” $\text{unit} \ \& \ co_X \mid \varepsilon \mid \varepsilon \cup \{rd_X, wr_X\}$ whenever $X \notin \text{locs}(\varepsilon)$. This means that the two computations are indistinguishable by environments that do not read, let alone modify X during the computation and assume regular read-write access once it is completed. It would alternatively be possible to replace the co -effect using a special set of private locations akin to the private regions from [10].

We use the notation $\text{rds}(\varepsilon)$, $\text{wrs}(\varepsilon)$, $\text{cos}(\varepsilon)$ to refer to the abstract locations l for which ε contains rd_l , wr_l , and co_l , respectively. We write $\text{locs}(\varepsilon) := \text{rds}(\varepsilon) \cup \text{wrs}(\varepsilon) \cup \text{cos}(\varepsilon)$. We also write ε^C for ε with all read effects removed and each wr_l in ε replaced by co_l .

Definition 5.1. An effect ε is well-formed (with respect to the current world) if $\text{locs}(\varepsilon) \subseteq \mathbf{w}$ and $\text{rds}(\varepsilon) \cap \text{cos}(\varepsilon) = \emptyset$ and $\text{cos}(\varepsilon) \subseteq \text{wrs}(\varepsilon)$. An effect specification is a triple $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ of well-formed effects such that $\varepsilon_2 \subseteq \varepsilon_3$.

An effect specification $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ approximates the behaviour of a computation e in the following way: the effect ε_1 summarizes side effects that may occur during the execution of e (corresponding to a guarantee condition in the rely-guarantee formalism [14]); the effect ε_2 summarizes effects of the interacting environment that e can tolerate while still functioning as expected (corresponding to a rely condition). Finally, ε_3 summarizes the side effects that may occur between start and completion of e . All the effects that the environment might introduce must be recorded in ε_3 because they are not under “our” control and might happen at any time even as the very last thing before the final result is returned. The effects flagged in ε_1 , on the other hand, do not necessarily show up in ε_3 , for a computation might be able to clean up those effects prior to returning the final result. The requirement that $\text{rds}(\varepsilon) \cap \text{cos}(\varepsilon) = \emptyset$ is owed to the fact that all effects should preserve their own precondition, however the precondition of rd_l is agreement on l which is not preserved by co_l . The requirement $\text{cos}(\varepsilon) \subseteq \text{wrs}(\varepsilon)$ reflects the fact that $\text{cos}(l)$ includes wr_l as a special case.

Note that if $\varepsilon^C \cup \varepsilon_1$ is a (well-formed) effect, then it is the case that $\text{rds}(\varepsilon_1) \cap (\text{wrs}(\varepsilon) \cup \text{cos}(\varepsilon)) = \emptyset$. We will use this observation to simplify some side conditions.

In our concrete examples, we abbreviate $\{co_l\} \cup \{wr_l\}$ by just co_l , in other words, the chaotic effect silently implies the write effect.

Consider the computations $e_1 = X := !X + 1; X := !X + 1$ and $e_2 = X := !X + 2$. Let ε_X stand for $\{rd_X, wr_X\}$ and analogously ε_Y . Each of the two computations can be assigned the effect $(\varepsilon_X, \varepsilon_Y, \varepsilon_X \cup \varepsilon_Y)$, but they are distinguishable at that effect typing. Under the looser specification $(\{co_{\varepsilon_X}\}, \varepsilon_Y, \varepsilon_X \cup \varepsilon_Y)$, however, they are indistinguishable, and our semantics is able to validate this equivalence, see Example 7.5.

Finally, consider the program $e = !X$ that simply reads a location storing an integer. We can show that this program has type $\mathbb{Z} \ \& \ \emptyset \mid \varepsilon \mid \varepsilon, rd_X$, where the read effect on X is only in the global effects.

Notations. For any well-formed effects $\varepsilon, \varepsilon'$ we use the notation $\varepsilon \perp \varepsilon'$ to mean that $\text{rds}(\varepsilon) \cap \text{wrs}(\varepsilon') = \text{rds}(\varepsilon') \cap \text{wrs}(\varepsilon) = \text{wrs}(\varepsilon) \cap \text{wrs}(\varepsilon') = \emptyset$. Note that this implies in particular $\text{cos}(\varepsilon) \cap \text{rds}(\varepsilon') = \emptyset$, etc. Intuitively, two programs exhibiting effects ε and ε' , respectively, commute with each other. We write $h \stackrel{\text{rds}(\varepsilon)}{\sim} h'$ to mean $h \stackrel{\sim}{\sim} h'$ for each $l \in \text{rds}(\varepsilon)$. We write $\stackrel{\varepsilon}{\rightarrow}$ for the transitive closure of $\bigcup_{l \in \text{wrs}(\varepsilon)} \stackrel{l}{\rightarrow} \cup \bigcup_{l \in \text{wrs}(\varepsilon)} \stackrel{l}{\rightarrow} \cap \stackrel{\sim}{\sim}$. Thus, $\stackrel{\varepsilon}{\rightarrow}$ allows steps by locations recorded as writing in ε and silent steps by all locations in the current world.

We define the notation $\varepsilon_1 \sqcup \varepsilon_2$ which appears in the parallel congruence rule by

$$\varepsilon_1 \sqcup \varepsilon_2 = \varepsilon_1 \cup \varepsilon_2 \setminus \{wr_l \mid wr_l \notin \varepsilon_1 \cap \varepsilon_2\} \setminus \{co_l \mid co_l \notin \varepsilon_1 \cap \varepsilon_2\}$$

6. Typing and congruence rules

Types are given by the grammar

$$\tau ::= \text{unit} \mid \text{int} \mid \text{bool} \mid A \mid \tau_1 \times \tau_2 \mid \tau_1 \xrightarrow[\varepsilon_2]{\varepsilon_1 \mid \varepsilon_3} \tau_2$$

where A ranges over user-specified abstract types. They will typically include reference types such as intref and also types like lists, sets, and even objects. In $\tau_1 \xrightarrow[\varepsilon_2]{\varepsilon_1 \mid \varepsilon_3} \tau_2$ the triple of effects $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ must be an effect specification.

We use two judgments:

- $\Gamma \vdash v \leq v' : \tau$ specifying that values v and v' have type τ and that v approximates v' ,
- $\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3$ specifying that the programs e and e' under the context Γ have type τ , with the effect specification $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ specifying, respectively, the effects during execution, the effects of the interacting environment and the start and completion effects. Moreover, e approximates e' at this specification.

We assume an ambient set of *axioms* each having the form (v, v', τ) where v, v' are values in the metalanguage and τ is a type meaning that v and v' are claimed to be of type τ and that v approximates v' . This must then be proved “manually” using the semantics rather than using the rules. The

We also define typing judgements $\Gamma \vdash v : \tau$ and $\Gamma \vdash e : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3$ which denote the special case when $\Gamma \vdash v \leq v : \tau$ and $\Gamma \vdash e \leq e : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3$ can be derived from the rules from Figure 6. We do not formulate explicit typing rules to save space.

The plan is to justify all the rules semantically using a logical relation (Section 7) and to then conclude their soundness w.r.t. typed observational approximation and equivalence (Section 8).

The parallel composition rule states that two programs e_1 and e_2 can be composed when their internal effects are not conflicting in the sense that the internal effects of one program appear as environment interaction effects of the other program. Note the relationship to the parallel composition rule of the rely-guarantee formalism [14]. Also note that the effects of computations e_1 and e_2 are not required to be independent from each other as we do in the parallelization rule further down.

The appearance of the \sqcup -operation deserves special mention. It might be, for example, that e_1 modifies X on the way, thus $wr_X \in \varepsilon_1$ but cleans up this modification by eventually restoring the old value of X . This would be reflected by $wr_X \notin \varepsilon \cup \varepsilon' \cup \varepsilon_2$. In that case, we would not expect to see wr_X in the end-to-end effect of the parallel composition and that is precisely what \sqcup achieves.

The rules labelled (Sem) make available all kinds of program transformations that are valid on the level of the *untyped* denotational semantics, including commuting conversions for let and if, fixpoint unrolling, and beta and eta equalities.

Finally, we have several effect-dependent (in)equalities: the parallelization rule generalises a similar rule from [10]. The other ones are concurrent version of analogous rules for sequential computation that have been analysed in previous work [6–8, 25] and are at the basis of all kinds of compiler optimizations. The side conditions on the effects are rather subtle and much less obvious than those found in a sequential setting. The parallelization rule is similar to the parallel congruence rule in that it requires the participating computations to mutually tolerate each other. This time, however, since the two computations being compared will do rather different things temporarily they must be oblivious against chaotic access, hence the $(-)^C$ strengthenings in the premise.

The reason for the appearance of $(-)^C$ in the other rules is similar. The rule for pure lambda hoist seems unusual and will thus be explained in more detail. First, the computation e_1 to be

hoisted may indeed have side effects ε_1 so long as they are cleaned up by the time e_1 completes and the intervening environment does not notice (modelled by the conditions $\varepsilon_1 \perp \varepsilon$ and final effect $\varepsilon^C = \varepsilon^C \cup \emptyset$). In the conclusion the transient effect ε_1 shows up again, but $(-)^C$ -ed since it only appears in different sides. Also in the other rules like commuting etc. it is the case that the familiar side conditions on applicability only affect the end-to-end effects whereas the transient effects are merely required not to interfere with the environment.

The following definitions provide the semantics of our effect annotations.

Definition 6.1 (Tiling). *Let $w \vdash \varepsilon$. We write $[\varepsilon](h, h', h_1, h'_1)$ to mean that (i) $h \models w \Rightarrow h \xrightarrow{\varepsilon} h_1$ and (ii) $h' \models w \Rightarrow h' \xrightarrow{\varepsilon} h'_1$ and (iii) $h \xrightarrow{\text{rds}(\varepsilon)} h'$ and $l \in \text{wrs}(\varepsilon) \setminus \text{cos}(\varepsilon)$ imply $(h \stackrel{l}{=} h_1 \wedge h' \stackrel{l}{=} h'_1) \vee h_1 \stackrel{l}{=} h'_1$.*

Thus, assuming semantic consistency of heaps, h and h' evolve to h_1 and h'_1 according to the modifying (writing or chaotic) locations in ε , and if h, h' agree on the reads of ε then written locations will either be identically modified or left alone.

If the step relations of all abstract locations commute with each other then tiling admits an alternative characterisation in terms of preservation of binary relations [8]. The present more operational version is inspired by the treatment of effects in [10].

Lemma 6.2. *Suppose that $w \vdash \varepsilon$, $w \vdash \varepsilon_1$, $w \vdash \varepsilon_2$. The following hold whenever well-formed.*

1. *If $[\varepsilon](h, h', h_1, h'_1)$ and $[\varepsilon](h_1, h'_1, h_2, h'_2)$ then $[\varepsilon](h, h', h_2, h'_2)$;*
2. *$[\varepsilon](h, h', h, h')$*
3. *If $\varepsilon_1 \subseteq \varepsilon_2$ then $[\varepsilon_1](h, h', h_1, h'_1) \Rightarrow [\varepsilon_2](h, h', h_1, h'_1)$*
4. *$[\varepsilon](h, h', h_1, h'_1) \Rightarrow [\varepsilon^C](h, h', h_1, h'_1)$*
5. *If $[\varepsilon](h, h', k, k')$ and $h \xrightarrow{\text{rds}(\varepsilon)} h'$ then $k \xrightarrow{\text{rds}(\varepsilon)} k'$. (this relies on $\text{rds}(\varepsilon) \cap \text{cos}(\varepsilon) = \emptyset$.)*
6. *Suppose $[\varepsilon](h, h', h_1, h'_1)$. If $h \models w$ then $h_1 \models w$; if $h' \models w$ then $h'_1 \models w$.*

7. Logical Relation

Definition 7.1 (Specifications). *A value specification is a relation $E \subseteq \mathbb{V} \times \mathbb{V}$ such that*

- *if $x_1 \leq x$ and $y \leq y_1$ and $x E y$ then $x_1 E y_1$;*
- *if $(x_i)_i$ and $(y_i)_i$ are chains such that $x_i E y_i$ then $\sup_i x_i E \sup_i y_i$, i.e., E is admissible qua relation;*
- *if $x E y$ then $p_i(x) E p_i(y)$ for each i , i.e. E is closed under the canonical deflations.*

Similarly, a computation specification is a relation $Q \subseteq T\mathbb{V} \times T\mathbb{V}$ such that $\leq; Q; \leq \subseteq Q$ and Q is admissible qua relation and Q is closed under the canonical deflations q_i .

The requirement $\leq; E; \leq \subseteq E$ ensures smooth interaction with the down-closure built into our trace monad. Admissibility is needed for the soundness of recursion and closure under the canonical deflations, finally is needed so that Lemma 3.3 can be applied.

Definition 7.2. *If $E \subseteq \mathbb{V} \times \mathbb{V}$ and $Q \subseteq T\mathbb{V} \times T\mathbb{V}$ then the relation $E \rightarrow Q \subseteq \mathbb{V} \times \mathbb{V}$ is defined by*

$$fE \rightarrow Qf' \iff \forall x x'. (x E x' \Rightarrow (f(x) Q f'(x')))$$

In particular, for $fE \rightarrow Qf'$ to hold, both f, f' must be functions (and not elements of base type or tuples).

Lemma 7.3. *If E and Q are specifications so is $E \rightarrow Q$.*

The following is the crucial definition of this paper; it gives a semantic counterpart to observational approximation and, due to its game-theoretic flavour, allows for very intuitive proofs.

Definition 7.4. *Let $E \subseteq \mathbb{V} \times \mathbb{V}$ be a value specification and $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ an effect specification. We define the relations $T_0(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ and $T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ between sets of trace-value pairs, i.e. on $\mathcal{P}(\text{Tr} \times \text{Values})$:*

$(U, U') \in T_0(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ if and only if

$$\left[\begin{array}{l} \forall ((h_1, k_1) \dots (h_n, k_n), a) \in U. h_1 \models w \Rightarrow \\ \forall h'_1. h'_1 \models w \Rightarrow h_1 \xrightarrow{\text{rds}(\varepsilon_3)} h'_1 \Rightarrow \\ \exists k'_1. [\varepsilon_1](h_1, h'_1, k_1, k'_1) \wedge \forall h'_2. [\varepsilon_2](k_1, k'_1, h_2, h'_2) \Rightarrow \\ \exists k'_2. [\varepsilon_1](h_2, h'_2, k_2, k'_2) \wedge \forall h'_3. [\varepsilon_2](k_2, k'_2, h_3, h'_3) \Rightarrow \\ \dots \\ \exists k'_n. [\varepsilon_1](h_n, h'_n, k'_n, k'_n) \wedge [\varepsilon_3](h_1, h'_1, k_n, k'_n) \wedge \\ \exists a' \in \mathbb{V}. (a, a') \in E \wedge ((h'_1, k'_1) \dots (h'_n, k'_n), a') \in U' \end{array} \right]$$

We define the relation $T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3) \subseteq T\mathbb{V} \times T\mathbb{V}$ as the admissible closure of T_0 , i.e. $\text{Adm}(T_0(E, \varepsilon_1, \varepsilon_2, \varepsilon_3))$.

The game-theoretic view of $T_0(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ may be understood as follows. Given $U, U' \in T\mathbb{V}$ we can consider a game between a proponent (who believes $(U, U') \in T\mathbb{V}$) and an opponent who believes otherwise. The game begins by the opponent selecting an element $((h_1, k_1) \dots (h_n, k_n), a) \in U$ and $h_1 \models w$, the *pilot trace* and a start heap $h'_1 \models w$ such that $h_1 \xrightarrow{\text{rds}(\varepsilon_3)} h'_1$ to begin a trace in U' . Then, the proponent answers with a matching heap k'_1 so that $[\varepsilon_1](h_1, h'_1, k_1, k'_1)$. If $h_1 \xrightarrow{\text{rds}(\varepsilon_1)} h'_1$ does not hold, proponent does not need to ensure that writes are in sync. The opponent then plays a heap h'_2 so that $[\varepsilon_2](k_1, k'_1, h_2, h'_2)$. At this point, it is in the proponent's interest to make sure that $k_1 \xrightarrow{\text{rds}(\varepsilon_2)} k'_1$ for otherwise opponent may make “funny” moves.

Then, again, proponent plays a heap k'_2 such that $[\varepsilon_1](h_2, h'_2, k_2, k'_2)$ and so on until, proponent has played k'_n so that $[\varepsilon_1](h_n, h'_n, k_n, k'_n)$. After that final heap has been played, it is checked that $[\varepsilon_3](h, h', k_n, k'_n)$ holds. If not, proponent loses. If yes, then proponent must also play a value a' and it is then checked whether or not $((h'_1, k'_1) \dots (h'_n, k'_n), a') \in U'$ and $(a E a')$. If this is the case or if at any one point in the game the opponent was unable to move because there exists no appropriate heap then the proponent has won the game. Otherwise the opponent wins and we have $(U, U') \in T_0(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ iff the proponent has a winning strategy for that game.

We notice that by Lemma 6.2(6) well-formedness of heaps w.r.t. the ambient world is a global invariant which allows us to refrain from explicitly assuming and asserting it in subsequent proofs and statements.

We now illustrate the game with a few examples.

Example 7.5. Consider the following programs: $e_1 = (X := !X + 1; X := !X + 1)$ and $e_2 = (X := !X + 2)$. Let $l = \text{int}(X)$ be the abstract location for a single integer stored at X (see Section 4). Let $E = \llbracket \text{unit} \rrbracket = \{(((), ()))\}$ be the value specification for the unit type.

We show that $(\llbracket e_1 \rrbracket, \llbracket e_2 \rrbracket) \in T(E, \{co_l\}, \varepsilon, \varepsilon \cup \{rd_l, wr_l\})$ under the assumption that $\{co_l\} \perp \varepsilon$, that is, when the environment does not read nor write X . This condition is clearly necessary, for e_1 and e_2 can be distinguished by an environment allowed to read or write X .

Let us now prove the claim when $\{co_l\} \perp \varepsilon$. The opponent picks a pilot trace in the semantics of e_1 , for example, $((h_1, k_1)(h_2, k_2), ())$ where $h_1(X) = n$ and $k_1(X) = n + 1$ and $h_2(X) = n'$ and $k_2(X) = n' + 1$. The other possible traces are stuttering or mumbling variants of this one and do not present additional difficulties. The opponent also chooses a heap h'_1 such that $h_1 \stackrel{l}{=} h'_1$, i.e., $h'_1(X) = n$. Now the proponent will choose to stutter for the time being and thus selects $k'_1 := h'_1$. Indeed, $[co_l](h_1, h'_1, k_1, k'_1)$ holds, so this is legal. The opponent now presents h'_2 such that $[\varepsilon](k_1, k'_1, h_2, h'_2)$. By the assumption on ε we know that $n' = h_2(X) = k_1(X) = n + 1$ and also $h'_2(X) = k'_1(X) = n$. The proponent now answers with

$$\begin{array}{c}
\frac{}{\Gamma \vdash \text{true} \leq \text{true} : \text{bool}} \quad \frac{}{\Gamma \vdash \text{false} \leq \text{false} : \text{bool}} \quad \frac{}{\Gamma \vdash n \leq n : \text{int}} \quad \frac{}{\Gamma, x : \tau \vdash x \leq x : \tau} \quad \frac{\Gamma \vdash v \leq v' : \tau_1 \times \tau_2}{\Gamma \vdash v.i \leq v'.i : \tau_i} \\
\frac{\Gamma \vdash e_1 \leq e_2 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3 \quad \Gamma \vdash e_1 \leq e_2 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3}{\Gamma \vdash e_1 \leq e_3 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3} \quad \frac{\Gamma \vdash v \leq v' : \tau}{\Gamma \vdash v \leq v' : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3} \quad \frac{\Gamma \vdash v_i \leq v'_i : \tau_1 \ i = 1, 2}{\Gamma \vdash (v_1, v_2) \leq (v'_1, v'_2) : \tau_1 \times \tau_2} \\
\frac{\Gamma \vdash v_1 \leq v'_1 : \tau_1 \xrightarrow{\varepsilon_1 \mid \varepsilon_3}{\varepsilon_2} \tau_2 \quad \Gamma \vdash v_2 \leq v'_2 : \tau_1}{\Gamma \vdash v_1 v_2 \leq v'_1 v'_2 : \tau_2 \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3} \quad \frac{\Gamma \vdash v \leq v' : \text{bool} \quad \Gamma \vdash e_1 \leq e'_1 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3 \quad \Gamma \vdash e_2 \leq e'_2 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3}{\Gamma \vdash \text{if } v \text{ then } e_1 \text{ else } e_2 \leq \text{if } v' \text{ then } e'_1 \text{ else } e'_2 : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3} \\
\frac{\Gamma \vdash e_1 \leq e'_1 : \tau_1 \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3 \quad \Gamma, x : \tau_1 \vdash e_2 \leq e'_2 : \tau_2 \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3}{\Gamma \vdash \text{let } x = e_1 \text{ in } e_2 \leq \text{let } x = e'_1 \text{ in } e'_2 : \tau_2 \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3} \quad \frac{\Gamma, f : \tau_1 \xrightarrow{\varepsilon_1 \mid \varepsilon_3}{\varepsilon_2} \tau_2, x : \tau_1 \vdash e \leq e' : \tau_2 \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3}{\Gamma \vdash \text{rec } f \ x = e \leq \text{rec } f \ x = e' : \tau_1 \xrightarrow{\varepsilon_1 \mid \varepsilon_3}{\varepsilon_2} \tau_2} \\
\frac{\Gamma \vdash e_1 \leq e'_1 : \tau_1 \ \& \ \varepsilon_1 \mid \varepsilon \cup \varepsilon_2 \mid \varepsilon \cup \varepsilon_2 \cup \varepsilon' \quad \Gamma \vdash e_2 \leq e'_2 : \tau_2 \ \& \ \varepsilon_2 \mid \varepsilon \cup \varepsilon_1 \mid \varepsilon \cup \varepsilon_1 \cup \varepsilon'}{\Gamma \vdash e_1 \parallel e_2 \leq e'_1 \parallel e'_2 : \tau_1 \times \tau_2 \ \& \ \varepsilon_1 \cup \varepsilon_2 \mid \varepsilon \mid \varepsilon \cup \varepsilon' \cup (\varepsilon_1 \cup \varepsilon_2)} \\
\frac{\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3 \quad \llbracket e \rrbracket = \llbracket e' \rrbracket}{\Gamma \vdash e' \leq e' : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3} \text{Sem}_1 \quad \frac{\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3 \quad \llbracket e \rrbracket = \llbracket e' \rrbracket}{\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3} \text{Sem}_2 \quad \frac{(v, v', \tau) \text{ an axiom}}{\Gamma \vdash v \leq v' : \tau} \text{Ax}_1 \\
\frac{\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3 \quad \varepsilon_1 \subseteq \varepsilon'_1 \quad \varepsilon'_2 \subseteq \varepsilon_2 \quad \varepsilon_3 \subseteq \varepsilon'_3}{\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon'_1 \mid \varepsilon'_2 \mid \varepsilon'_3} \quad \frac{\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon_1 \mid \emptyset \mid \varepsilon_3}{\Gamma \vdash \text{atomic}(e) \leq \text{atomic}(e') : \tau \ \& \ \varepsilon_3 \mid \varepsilon_2 \mid \varepsilon_2 \cup \varepsilon_3} \text{Atom} \quad \frac{(v, v', \tau) \text{ an axiom}}{\Gamma \vdash v' \leq v' : \tau} \text{Ax}_2
\end{array}$$

Figure 4. Typing and congruence rules

$$\begin{array}{c}
\frac{\Gamma \vdash e_1 : \tau_1 \ \& \ \varepsilon_1 \mid \varepsilon^C \cup \varepsilon_2^C \mid \varepsilon^C \cup \varepsilon_2^C \cup \varepsilon'_1 \quad \Gamma \vdash e_2 : \tau_2 \ \& \ \varepsilon_2 \mid \varepsilon^C \cup \varepsilon_1^C \mid \varepsilon^C \cup \varepsilon_1^C \cup \varepsilon'_2 \quad \varepsilon_1 \perp \varepsilon_2 \quad \varepsilon_1 \perp \varepsilon \quad \varepsilon_2 \perp \varepsilon}{\Gamma \vdash e_1 \parallel e_2 \leq (\text{let } x = e_1 \text{ in let } y = e_2 \text{ in } (x, y)) : \tau_1 \times \tau_2 \ \& \ \varepsilon_1^C \cup \varepsilon_2^C \mid \varepsilon \mid \varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2} \text{Parallelization} \\
\frac{\Gamma \vdash e_1 : \tau_1 \ \& \ \varepsilon_1 \mid \varepsilon^C \mid \varepsilon^C \cup \varepsilon'_1 \quad \Gamma \vdash e_2 : \tau_2 \ \& \ \varepsilon_2 \mid \varepsilon^C \mid \varepsilon^C \cup \varepsilon'_2 \quad \varepsilon'_1 \perp \varepsilon'_2 \quad \varepsilon_1 \perp \varepsilon \quad \varepsilon_2 \perp \varepsilon}{\Gamma \vdash (\text{let } x = e_1 \text{ in let } y = e_2 \text{ in } (x, y)) = (\text{let } y = e_2 \text{ in let } x = e_1 \text{ in } (x, y)) : \tau_1 \times \tau_2 \ \& \ \varepsilon_1^C \cup \varepsilon_2^C \mid \varepsilon \mid \varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2} \text{Commuting} \\
\frac{\Gamma \vdash e : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2^C \mid \varepsilon_2^C \cup \varepsilon' \quad \text{rds}(\varepsilon') \cap \text{wrs}(\varepsilon') = \emptyset \quad \varepsilon_2 \perp \varepsilon_1}{\Gamma \vdash (\text{let } x = e \text{ in } (x, x)) \leq (\text{let } x = e \text{ in let } y = e \text{ in } (x, y)) : \tau \times \tau \ \& \ \varepsilon_1^C \mid \varepsilon_2 \mid \varepsilon_2 \cup \varepsilon'} \text{Duplicated} \\
\frac{(v, v', \tau) \text{ an axiom} \quad \Gamma \vdash v \leq v' : \tau \quad \text{Ax} \quad \Gamma \vdash e_1 : \tau_1 \ \& \ \varepsilon_1 \mid \varepsilon^C \mid \varepsilon^C \quad \Gamma, x : \tau_3, y : \tau_1 \vdash e_2 : \tau_2 \ \& \ \varepsilon_2 \mid \varepsilon \mid \varepsilon \cup \varepsilon_2 \quad \varepsilon \perp \varepsilon_1}{\Gamma \vdash \text{let } y = e_1 \text{ in } \lambda x. e_2 \leq \lambda x. \text{let } y = e_1 \text{ in } e_2 : \tau_3 \xrightarrow[\varepsilon]{\varepsilon_1^C \cup \varepsilon_2 \mid \varepsilon \cup \varepsilon_3} \tau_2 \ \& \ \varepsilon_1^C \mid \varepsilon \mid \varepsilon} \text{Lambda Hoist} \\
\frac{\Gamma \vdash e_1 : \tau_1 \ \& \ \varepsilon_1 \mid \varepsilon^C \mid \varepsilon^C \cup \varepsilon'_1 \quad \Gamma \vdash e_2 : \tau_2 \ \& \ \varepsilon_2 \mid \varepsilon \mid \varepsilon'_2 \quad \varepsilon_1 \perp \varepsilon \quad \text{wrs}(\varepsilon'_1) = \emptyset}{\Gamma \vdash e_2 \leq (\text{let } x = e_1 \text{ in } e_2) : \tau_2 \ \& \ \varepsilon_1^C \cup \varepsilon_2 \mid \varepsilon \mid \varepsilon \cup \varepsilon'_2} \text{Deadcode}
\end{array}$$

Figure 5. Effect-dependent transformations.

$k'_2 := h'_2[X \mapsto n + 2]$. It follows that $[co_1](h_2, h'_2, k_2, k'_2)$ and also $[rd_1, wr_1](h_1, h'_1, k_2, k'_2)$. Finally, by stuttering $(h'_1, h'_1)(h'_2, h'_2[X \mapsto n + 2]) \in \llbracket e_2 \rrbracket$ so that proponent wins the game.

Example 7.6. Consider the following programs e_1 and e_2 :

$(X := !X + 1 \parallel Y := !Y + 1)$ and $(X := !X + 1; Y := !Y + 1)$. We show $(\llbracket e_1 \rrbracket, \llbracket e_2 \rrbracket) \in T(E, \{co_X, co_Y\}, \varepsilon, \varepsilon \cup \{rd_X, rd_Y, wr_X, wr_Y\})$, provided ε does not read nor modify X and Y . This equivalence could be deduced syntactically using our parallelization equation shown in Figure 5. For illustrative purpose, however, we describe its semantic proof using a game.

The opponent picks a pilot trace in $\llbracket e_1 \rrbracket$, for example, the trace $([n_1|n_2], [n_1|n_2 + 1])([n_1|n_2 + 1], [n_1 + 1|n_2 + 1])((), ())$, where $[n_X|n_Y]$ denotes a heap where X and Y store n_X and n_Y , respectively. Notice that in this trace, Y is incremented before X and since ε does not read nor modify X and Y , the environment move does not change the values in X nor Y . We are also given an initial heap

h'_1 that agrees with the initial heap $[n_1|n_2]$ on the reads of $\varepsilon \cup \{rd_X, rd_Y, wr_X, wr_Y\}$. Thus, h'_1 should be of the form $[n_1|n_2]$.

We now play the move $([n_1|n_2], [n_1 + 1|n_2])$. This is a valid move in the game as $[co_X, co_Y]([n_1|n_2], [n_1|n_2], [n_1|n_2 + 1], [n_1 + 1|n_2])$. The environment moves returning $[n_1 + 1|n_2]$ as it does not read nor modify X and Y . We can now match the trace above by playing $([n_1 + 1|n_2], [n_1 + 1|n_2 + 1])$ and returning $((), ())$, winning the game.

The following is one of the main technical result of our paper and shows that the computation specifications $T(\dots)$ can indeed serve as the basis for a logical relation. We just show here the soundness proof for the parallel congruence rule. The missing proofs appear in the attached Appendix.

Theorem 7.7. *The following hold whenever well-formed.*

1. If $(U, U') \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ then $(q_i(U), q_i(U')) \in T(E, \varepsilon_1, \varepsilon_2)$.
2. $T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ is a computation specification.
3. If $(U, U') \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ then $(U^\dagger, U'^\dagger) \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$.

4. If $(a, a') \in E$ then $(\text{rtn}(a), \text{rtn}(a'))$ is in $T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$.
5. Suppose that $(\varepsilon_1, \varepsilon_2, \varepsilon_3)$ is an effect specification where $\varepsilon_1 \cup \varepsilon_2 \subseteq \varepsilon_3$. Suppose that whenever $h \xrightarrow{\text{rds}(\varepsilon_1)} h'$ and $c(h) = (h_1, a)$ then there exist (h'_1, a') such that $c'(h') = (h'_1, a')$ and $[\varepsilon_1](h, h', h_1, h'_1)$ and aEa' . We then have for any ε_2 , $(\text{fromstate}(c), \text{fromstate}(c')) \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$.
6. If $(f, f') \in E_1 \rightarrow T(E_2, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ and $(U, U') \in T(E_1, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ then $(\text{bnd}(f, U), \text{bnd}(f', U')) \in T(E_2, \varepsilon_1, \varepsilon_2, \varepsilon_3)$.
7. If $(U_1, U'_1) \in T(E_1, \varepsilon_1, \varepsilon \cup \varepsilon_2, \varepsilon \cup \varepsilon_1 \cup \varepsilon')$ and $(U_2, U'_2) \in T(E_2, \varepsilon_2, \varepsilon \cup \varepsilon_1, \varepsilon \cup \varepsilon_1 \cup \varepsilon')$ then $(U_1 \mid U'_1, U_2 \mid U'_2) \in T(E_1 \times E_2, \varepsilon_1 \cup \varepsilon_2, \varepsilon, \varepsilon \cup \varepsilon' \cup (\varepsilon_1 \sqcup \varepsilon_2))$.
8. $(U, U') \in T(E, \varepsilon_1, \emptyset, \varepsilon_3) \Rightarrow (at(U), at(U')) \in T(\varepsilon_3, \varepsilon_2, \varepsilon_2 \cup \varepsilon_3)$.

Proof. Ad 7. Suppose that $(U_1, U'_1) \in T(E_1, \varepsilon_1, \varepsilon \cup \varepsilon_2, \varepsilon \cup \varepsilon_1 \cup \varepsilon')$ and $(U_2, U'_2) \in T(E_2, \varepsilon_2, \varepsilon \cup \varepsilon_1, \varepsilon \cup \varepsilon_1 \cup \varepsilon')$ and let $(t, (a, b)) \in U_1 \mid U_2$, thus $\text{inter}(t_1, t_2, t)$ (ignoring \dagger by item 3) where $(t_1, a) \in U_1$ and $(t_2, b) \in U_2$. Let S_1, S_2 be corresponding winning strategies. The idea is to use S_1 when we are in t_1 and to use S_2 when we are in t_2 . Supposing that t starts with a t_1 fragment we begin by playing according to S_1 . Let t be of the form:

$$t = (h_1, k_1) \cdots (h_n, k_n)(h_{n+1}, k_{n+1}) \cdots (h_{n+m}, k_{n+m}) \\ (h_{n+m+1}, k_{n+m+1}) \cdots (h_{n+m+k}, k_{n+m+k}) \cdots (h_p, k_p)$$

composed of pieces of the traces t_1 and t_2 . Assume w.l.o.g. that the first piece $(h_1, k_1) \cdots (h_n, k_n)$ is a part of t_1 . We are given a initial heap h'_1 such that $h \xrightarrow{\text{rds}(\varepsilon \cup \varepsilon' \cup (\varepsilon_1 \sqcup \varepsilon_2))} h'$. Since $\text{rds}(\varepsilon_1 \sqcup \varepsilon_2) = \text{rds}(\varepsilon_1) \cup \text{rds}(\varepsilon_2)$, we can apply strategy S_1 to guide us through the first part of the game, obtaining:

$$(h'_1, k'_1) \cdots (h'_n, k'_n)$$

Moreover, we have an environment move which forms the tile $[\varepsilon](k_n, k'_n, h_{n+1}, h'_{n+1})$. Thus, we have the tile $[\varepsilon \cup \varepsilon_1](h_1, h'_1, h_{n+1}, h'_{n+1})$ which can be seen as an environment move for t_2 . Therefore, we can use strategy S_2 for the U' and continue the game, obtaining the trace piece:

$$(h'_{n+1}, k'_{n+1}) \cdots (h'_{n+m}, k'_{n+m})$$

Now, we can return to the S_1 game as the trace above is seen as an environment move for U . Alternating these strategies, we get a trace t which is in $(U \mid U')$. Let (a', b') be the final values reached at the end. It is clear that $[\varepsilon \cup \varepsilon' \cup \varepsilon_1 \cup \varepsilon_2](h, h', h_p, h'_p)$ and also aE_1a' and bE_2b' .

It remains to assert the stronger statement $[\varepsilon \cup \varepsilon' \cup (\varepsilon_1 \sqcup \varepsilon_2)](h, h', h_p, h'_p)$. To see this suppose that $wr_1 \in \varepsilon_1 \setminus \varepsilon_2 \setminus \varepsilon \setminus \varepsilon'$. Since the entire game can be viewed as an instance of the game U_1 vs U'_1 with interventions by U_2 vs. U'_2 regarded as environmental interactions we have $[\varepsilon \cup \varepsilon_2 \cup \varepsilon'](h, h', h_p, h'_p)$ so that in fact $h \stackrel{!}{=} h_p$ and $h' \stackrel{!}{=} h'_p$. The case of co_1 and $\varepsilon_1, \varepsilon_2$ interchanged is analogous. \square

We assign a value specification $\llbracket \tau \rrbracket$ to each refined type by

- $\llbracket \text{int} \rrbracket = \{(v, v') \mid v = v' \in \mathbb{Z}\}$ • $\llbracket \tau_1 \times \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket \times \llbracket \tau_2 \rrbracket$
- $\llbracket \tau_1 \xrightarrow{\varepsilon_1 \mid \varepsilon_2} \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket \rightarrow T(\llbracket \tau_2 \rrbracket, \varepsilon_1, \varepsilon_2, \varepsilon_3)$

We omit the obvious definition of the other basic types and assume value specifications for user-specified types as given.

Assumption 1. We henceforth adopt the following soundness assumption which must be established concretely for every concrete instance of our framework.

- The initial heap satisfies the current world: $h_{\text{init}} \models w$.
- Each axiom is type sound: whenever (v, v', τ) is an axiom then $(v, v) \in \llbracket \tau \rrbracket$ and $(v', v') \in \llbracket \tau \rrbracket$.
- Each axiom is inequationally sound: whenever (v, v', τ) is an axiom then $(v, v') \in \llbracket \tau \rrbracket$.

Theorem 7.8. Suppose that $\Gamma \vdash v : \tau$ and $\Gamma \vdash e : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3$. Then $(\eta, \eta') \in \llbracket \Gamma \rrbracket$ (interpreting a context as a cartesian product) implies $(\llbracket v \rrbracket \eta, \llbracket v \rrbracket \eta') \in \llbracket \tau \rrbracket$ and $(\llbracket e \rrbracket \eta, \llbracket e \rrbracket \eta') \in T(\llbracket \tau \rrbracket, \varepsilon_1, \varepsilon_2, \varepsilon_3)$.

Proof. By induction on derivations. Most cases are already subsumed by Theorem 7.7. The typing rules regarding functions and recursion follow from the definitions and from the fact that all specifications are admissible. \square

8. Typed observational approximation

Definition 8.1 (Observational approximation). Let v, v' be value expressions where $\vdash v : \tau$ and $\vdash v' : \tau$. We say that v observationally approximates v' at type τ if for all f such that $\vdash f : \tau \xrightarrow{\varepsilon_1 \mid \varepsilon_3} \text{int}$ (“observations”) it is the case that if $((h_{\text{init}}, k), n) \in \llbracket f \ v \rrbracket$ for $v \in \mathbb{Z}$ and starting from h_{init} then $((h_{\text{init}}, k'), n) \in \llbracket f \ v' \rrbracket$ for some k' . We write $\vdash v \leq_{\text{obs}} v'$ in this case. We say that v and v' are observationally equivalent at type τ , written $\vdash v =_{\text{obs}} v'$ if both $\vdash v \leq_{\text{obs}} v' : \tau$ and $\vdash v' \leq_{\text{obs}} v : \tau$.

This means that for every test harness f we build around v and v' , no matter how complicated it is and whatever environments it sets up to run concurrently with v and v' it is the case that each terminating computation of v (in the environment installed by f) can be matched by a terminating computation with the same result by v' in the same environment. It is important, however, that the environment be well typed, thus will respect the contracts set up by the type τ . E.g. if τ is a functional type expecting, say, a pure function as argument then, by the typing restriction, the environment f cannot suddenly feed v and v' a side-effecting function as input.

We remark that observational approximation extends canonically to open terms by lambda abstracting free variables (and adding a dummy abstraction in the case of closed terms) [6].

As usual, the logical relation is sound with respect to typed observational approximation and thus can be used to deduce nontrivial observational approximation relations. We state and prove the precise formulation of this result.

Theorem 8.2. Let v, v' be closed values and suppose that $(\llbracket v \rrbracket, \llbracket v' \rrbracket) \in \llbracket \tau \rrbracket$. Then $\vdash v \leq_{\text{obs}} v' : \tau$.

Proof. If $\vdash f : \tau \xrightarrow{\varepsilon_1 \mid \varepsilon_3} \text{int}$ then by Thm 7.8 we have $(\llbracket f \rrbracket, \llbracket f \rrbracket) \in \llbracket \tau \xrightarrow{\varepsilon_1 \mid \varepsilon_3} \text{int} \rrbracket$, so $(\llbracket f \ v \rrbracket, \llbracket f \ v' \rrbracket) \in T(\llbracket \text{int} \rrbracket, \varepsilon_1, \varepsilon_2, \varepsilon_3)^+$.

Let $((h_{\text{init}}, k), v) \in \llbracket f \ v \rrbracket$. We have $h_{\text{init}} \models w$ and thus in particular $h_{\text{init}} \xrightarrow{\text{rds}(\varepsilon_3) \cup \text{rds}(\varepsilon_1)} h_{\text{init}}$. There must therefore exist a matching heap k' and a value v' such that $((h_{\text{init}}, k'), v') \in \llbracket f \ v' \rrbracket$ and $v = v' \in \mathbb{Z}$. \square

This means that the examples from earlier on give rise to valid transformations in the sense of observational approximation. For instance, for e_1 and e_2 from Example 7.5 we find that $\lambda_{\dots} e_1 =_{\text{obs}} \lambda_{\dots} e_2$ at type unit $\xrightarrow{\{co_1\} \mid \varepsilon \cup \{rd_1, wr_1\}} \text{unit}$ whenever X does not appear in ε .

9. Effect-dependent transformations

We will now establish the semantic soundness of the inequational theory of effect-dependent program transformations given in Figure 5. It includes concurrent versions of the effect-dependent equations from [8, 25], but the side conditions on the environmental interaction are by no means obvious. We also note that some equations now only hold in one direction thus become inequations. This is in particular the case for duplicated computations. Suppose that $?$ is a computation that nondeterministically chooses a boolean value and let $e := \text{let } x = ? \text{ in } (x, x)$. Then, even though $?$ does

not read nor write any location we only have $e \leq (?, ?)$, but not $(?, ?) \leq e$ for $(?, ?)$ admits the result $(\text{true}, \text{false})$ but e does not. Furthermore, due to presence of nontermination the equations for dead code elimination and pure lambda hoist also hold in one direction only. It might be possible to restore both directions of said equations by introducing special effects for nondeterminism and nontermination; we have not explored this avenue. We concentrate the individual effect-dependent transformations before summarising the foregoing results in the general soundness Theorem 9.2.

In many of the equations, co-effects play an important role. For example, in the commuting and parallelization equations, the internal effects ε_1 and ε_2 in the premises are replaced by ε_1^C and ε_2^C in the internal effects of the conclusion. This makes sense intuitively because the computations are run in a different order, so for the internal moves, the locations in ε_1 and ε_2 can be modified in any way (see Example 7.6). However, in the global effect, we can still guarantee the effects ε'_1 and ε'_2 because of the \perp -conditions. This intuition appears directly in the soundness proofs.

The following thus constitutes the second main technical result of our paper. We sketch the soundness proof for parallelization. The detailed proofs appear in the attached Appendix.

Theorem 9.1. *The following hold whenever well-formed.*

- **Commuting** If $(U_1, U'_1) \in T(E_1, \varepsilon_1, \varepsilon_1^C, \varepsilon_1^C \cup \varepsilon'_1)$ and $(U_2, U'_2) \in T(E_2, \varepsilon_2, \varepsilon_2^C, \varepsilon_2^C \cup \varepsilon'_2)$ and $\varepsilon_1 \perp \varepsilon$ and $\varepsilon_2 \perp \varepsilon$ and $\varepsilon'_1 \perp \varepsilon'_2$ then
$$\begin{aligned} & ((t_1 t_2, (v_1, v_2)) \mid (t_1, v_1) \in U_1, (t_2, v_2) \in U_2)^\dagger, \\ & ((t'_1 t'_2, (v'_1, v'_2)) \mid (t'_1, v'_1) \in U'_1, (t'_2, v'_2) \in U'_2)^\dagger) \\ & \in T(E_1 \times E_2, (\varepsilon_1 \cup \varepsilon_2)^C, \varepsilon, \varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2) \end{aligned}$$
- **Duplicated** If $(U, U') \in T(E, \varepsilon_1, \varepsilon_2^C, \varepsilon_2^C \cup \varepsilon')$ and $\text{rds}(\varepsilon') \cap \text{wrs}(\varepsilon') = \emptyset$ and $\varepsilon_2 \perp \varepsilon_1$, then
$$\begin{aligned} & ((t, (v, v)) \mid (t, v) \in U)^\dagger, ((t'_1 t'_2, (v'_1, v'_2)) \mid (t'_1, v'_1) \in U', \\ & (t'_2, v'_2) \in U')^\dagger) \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_2 \cup \varepsilon') \end{aligned}$$
- **Pure** Let $(U, U') \in T(E, \varepsilon_1, \varepsilon_2^C, \varepsilon_2^C)$, such that $\varepsilon_1 \perp \varepsilon_2$. If $((q_1, k_1) \dots (q_n, k_n), v) \in U$ for some arbitrary trace $t = (q_1, k_1) \dots (q_n, k_n)$ (with $q_1 \models \mathbf{w}$) and value v , then $(\text{rtn}(v), U') \in T(E, \varepsilon_1^C, \varepsilon_2, \varepsilon_2)$;
- **Dead** Suppose that $(U, U') \in T(\text{unit}, \varepsilon_1, \varepsilon_2, \varepsilon_2 \cup \varepsilon'_1)$, where $\text{wrs}(\varepsilon'_1) = \emptyset$ and $\varepsilon_1 \perp \varepsilon_2$. Then $(U, \text{rtn}()) \in T(\text{unit}, \varepsilon_1^C, \varepsilon_2, \varepsilon_2 \cup \varepsilon'_1)$.
- **Parallelization** If $(U_1, U'_1) \in T(E_1, \varepsilon_1, \varepsilon_1^C \cup \varepsilon_2^C, \varepsilon_1^C \cup \varepsilon_2^C \cup \varepsilon'_1)$ and $(U_2, U'_2) \in T(E_2, \varepsilon_2, \varepsilon_2^C \cup \varepsilon_1^C, \varepsilon_2^C \cup \varepsilon_1^C \cup \varepsilon'_2)$ and $\varepsilon_1 \perp \varepsilon_2$ and $\varepsilon_1 \perp \varepsilon$ and $\varepsilon_2 \perp \varepsilon$, then
$$(U_1 \parallel U_2, \{(t'_1 t'_2, (v'_1, v'_2)) \mid (t'_1, v'_1) \in U'_1, (t'_2, v'_2) \in U'_2\}^\dagger) \in T(E_1 \times E_2, \varepsilon_1^C \cup \varepsilon_2^C, \varepsilon, \varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2)$$

Proof. (Sketch) **Parallelization.**

Assume w.l.o.g. that the pilot trace takes the form $(t, (v_1, v_2))$ where $\text{inter}(t_1, t_2, t)$ and $(t_i, v_i) \in U_i$. Just as in the commuting case we set up two side games U_i vs. U'_i on t_i, v_i . Unlike, in that case, however, these games are running simultaneously and along with the main game. Moves by the environment in the main game are forwarded to the side game we are currently in, i.e., the one to which the current portion of t being played on belongs. At each change of control, we switch between the two side games making last sequence of moves of the other game into a single environment move. It is here that the resilience against chaotic modification is needed. Once the play is over we then assert the claims about the end-to-end effect $\varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2$ location by location using the definition of tiling. \square

Theorem 9.2. *Suppose that $\Gamma \vdash v \leq v' : \tau$ and $\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3$ and assume that for each axiom (v, v', τ) it holds that $(v, v') \in \llbracket \tau \rrbracket^+$. Then $(\eta, \eta') \in \llbracket \Gamma \rrbracket^+$ (interpreting a*

context as a cartesian product) implies $(\llbracket v \rrbracket \eta, \llbracket v' \rrbracket \eta') \in \llbracket \tau \rrbracket^+$ and $(\llbracket e \rrbracket \eta, \llbracket e' \rrbracket \eta') \in T(\llbracket \tau \rrbracket, \varepsilon_1, \varepsilon_2, \varepsilon_3)^+$.

Sketch. In essence the proof is by induction on derivations of inequalities. However, we need to slightly strengthen the induction hypothesis as follows:

Define

$$\begin{aligned} \llbracket \Gamma \vdash \tau \rrbracket &= \{(f, f') \mid \forall (\eta, \eta') \in \llbracket \Gamma \rrbracket. (f(\eta), f'(\eta')) \in \llbracket \tau \rrbracket\} \\ \llbracket \Gamma \vdash \tau \&(\varepsilon_1, \varepsilon_2, \varepsilon_3) \rrbracket &= \{(f, f') \mid \forall (\eta, \eta') \in \llbracket \Gamma \rrbracket. \\ & (f(\eta), f'(\eta')) \in T(\llbracket \tau \rrbracket, \varepsilon_1, \varepsilon_2, \varepsilon_3)\} \end{aligned}$$

We now show by induction on derivations that $\Gamma \vdash v \leq v' : \tau$ implies $(\llbracket v \rrbracket, \llbracket v' \rrbracket) \in \llbracket \Gamma \vdash \tau \rrbracket^+$ and that $\Gamma \vdash e \leq e' : \tau \ \& \ \varepsilon_1 \mid \varepsilon_2 \mid \varepsilon_3$ implies $(\llbracket e \rrbracket, \llbracket e' \rrbracket) \in \llbracket \Gamma \vdash \tau \&(\varepsilon_1, \varepsilon_2, \varepsilon_3) \rrbracket^+$.

The various cases now follow from earlier results in a straightforward manner. Namely, we use Theorem 7.7 for the congruence rules and Theorem 9.1 for the effect-dependent transformations.

As a representative case we show the case where $e \equiv \text{let } x = e_1 \text{ in } e_2$ and $e' \equiv \text{let } x = e'_1 \text{ in } e'_2$. Inductively, we know $(\llbracket e_1 \rrbracket, \llbracket e'_1 \rrbracket) \in \llbracket \Gamma \vdash \tau_1 \&(\varepsilon_1, \varepsilon_2, \varepsilon_3) \rrbracket^{n_1}$ and $(\llbracket e_1 \rrbracket, \llbracket e'_1 \rrbracket) \in \llbracket \Gamma, x : \tau_1 \vdash \tau \&(\varepsilon_1, \varepsilon_2, \varepsilon_3) \rrbracket^{n_2}$ for some $n_1, n_2 > 0$. By Theorem 7.8, we also have $(\llbracket e_1 \rrbracket, \llbracket e'_1 \rrbracket) \in \llbracket \Gamma \vdash \tau_1 \&(\varepsilon_1, \varepsilon_2, \varepsilon_3) \rrbracket$ and analogous statements for e'_1, e_2, e'_2 . We can, therefore, assume, w.l.o.g. that $n_1 = n_2$ and then use Theorem 7.7 (6) repeatedly (n_1 times) so as to conclude $(\llbracket e \rrbracket, \llbracket e' \rrbracket) \in \llbracket \Gamma \vdash \tau \&(\varepsilon_1, \varepsilon_2, \varepsilon_3) \rrbracket^{n_1}$.

The rules for dead code and pure lambda hoist rely on the cases “Dead” and “Pure” of Thm 9.1 in a slightly indirect way. We sketch the argument for pure lambda hoist. The pilot trace begins with a trace belonging to e_1 and yielding a value v for x . We can then invoke case “Pure” on subsequent occurrences of e_1 in the right hand side. \square

Theorem 9.3. *Suppose that $\vdash v : \tau$ and $\vdash v' : \tau$ and that $(\llbracket v \rrbracket, \llbracket v' \rrbracket) \in \llbracket \tau \rrbracket^+$ where $(-)^+$ denotes transitive closure. Then $\vdash v \leq_{\text{obs}} v' : \tau$.*

Proof. If $\vdash f : \tau_1 \xrightarrow[\varepsilon_2]{\varepsilon_1 \mid \varepsilon_3} \text{int}$ then by Thm 7.8 we have $(\llbracket f \rrbracket, \llbracket f \rrbracket) \in$

$\llbracket \tau \xrightarrow[\varepsilon_2]{\varepsilon_1 \mid \varepsilon_3} \text{int} \rrbracket$, so $(\llbracket f v \rrbracket, \llbracket f v' \rrbracket) \in T(\llbracket \text{int} \rrbracket, \varepsilon_1, \varepsilon_2, \varepsilon_3)^+$.

Let $((h_{\text{init}}, k), v) \in \llbracket f v \rrbracket$. We have $h_{\text{init}} \models \mathbf{w}$ and thus in particular $h_{\text{init}} \xrightarrow[\text{rds}(\varepsilon_3) \cup \text{rds}(\varepsilon_1)]{\text{rds}(\varepsilon_3) \cup \text{rds}(\varepsilon_1)} h_{\text{init}}$. There must therefore exist a matching heap k' and a value v' such that $((h_{\text{init}}, k'), v') \in \llbracket f v' \rrbracket$ and $v = v' \in \mathbb{Z}$. \square

We now return to the examples that we discussed in Section 1 and demonstrate how to prove using our denotational semantics the properties that have been discussed informally.

Overlapping References With this example, we illustrate the parallelization rule. In particular, the functions declared in Section 1 have the following type, where ε does not read nor write X :

$$\begin{aligned} \text{readFst} : \text{unit} & \xrightarrow[\varepsilon^C, \text{coSnd}(X)]{\emptyset \mid \varepsilon^C, \text{coSnd}(X), \text{rdFst}(X)} \text{int} \\ \text{writeFst} : \text{int} & \xrightarrow[\varepsilon^C, \text{coSnd}(X)]{\text{wrFst}(X) \mid \varepsilon^C, \text{coSnd}(X), \text{wrFst}(X)} \text{unit} \end{aligned}$$

The obvious and analogous typings for readSnd and writeSnd are elided. We justify this typing semantically as described in Theorem 7.7. To illustrate how this is done, consider the function $(\text{writeSnd } 17)$. We show how the game is played against itself using the typing shown above. We start with a “pilot trace”, say: $([2|3], [2|3]), ([2|17], [2|17]), ()$ where $[x|y]$ denotes a store with $X = p(x, y)$ and other components left out for simplicity. The first step corresponds to our reading of X and in the second step – since there was no environment intervention – we write 17 into the first component.

We now start to play: Say that we start at the heap $[13|12]$. We answer $[13|12]$. If the environment does not change X , then

we write 17 to its first component resulting in the following trace, which is possible for `writeFst(17)`.

$(([13|12], [13|12]), ([13|12], [17|12]), (0))$

If, however, the environment plays [18|21] (a modification of both components of X has occurred), then we answer [17|21]. Again,

$(([13|12], [13|12]), ([18|21], [17|21]), (0))$

is a possible trace for `writeFst(17)`. It is easy to check that there is a strategy that justifies the typing given above.

Now, consider a program, e_1 , that only calls `readFst`, `writeFst`, and another program, e_2 , that only calls `readSnd`, `writeSnd`. Since the former functions have disjoint effects to the latter ones, e_1 and e_2 will have effect specifications, respectively, of the form $(\varepsilon_1, \varepsilon^C \cup \varepsilon_2^C, \varepsilon^C \cup \varepsilon_2^C \cup \varepsilon_1)$ and $(\varepsilon_2, \varepsilon^C \cup \varepsilon_1^C, \varepsilon^C \cup \varepsilon_1^C \cup \varepsilon_2)$, where $\varepsilon_1 \cap \varepsilon_2 = \varepsilon_1 \cap \varepsilon = \varepsilon_2 \cap \varepsilon = \emptyset$. Thus we can use the parallelization rule shown in Figure 5 to conclude that the behavior of $e_1 || e_2$ is the same as executing these programs sequentially, although they read and write to the same concrete location.

Loop Parallelization We show that the function map is equivalent to `map2Par`. It is easy to see that the function map is equivalent to the program `map2Seq`, which is the program obtained from `map2Par` by replacing the underlined parallel operator ‘ $||$ ’ in `map2Par` by a sequential operator ‘ $;$ ’. The proof goes simply by unfolding map.

We then proceed by showing `map2Seq` and `map2Par` are equivalent using our equations and the abstract locations `listodd(X)` and `listodd(X)` defined above. The piece of code that applies f first, namely $e_1 = n.\text{ele} := f(n.\text{ele})$, has global effects $\varepsilon'_1 = rd_{listodd(X)}, wr_{listodd(X)}$, while the second application, namely, $e_2 = n.\text{next.ele} := f(n.\text{next.ele})$, has effects $\varepsilon'_2 = rd_{listeven(X)}, wr_{listeven(X)}$. Notice that $\varepsilon'_1 \perp \varepsilon'_2$. Therefore, provided that the environment does not read nor modify the list, we can apply the parallelization equation to justify running e_1 and e_2 parallel is equivalent to running them in sequence.

Michael-Scott Queue We now show that the enqueue and dequeue functions described in Section 1 for the Michael-Scott Queue have the same behavior as their atomic versions. We only show the case for dequeue, as the case for enqueue is similar. More precisely, we now justify the axiom

$$(\text{dequeue}, \text{atomic}(\text{dequeue}), \text{unit} \xrightarrow[\text{MSQ}]{\text{MSQ}|\text{MSQ}} \text{int})$$

where $\text{MSQ} = \{rd_{\text{msq}(X)}, wr_{\text{msq}(X)}\}$. That is, they approximate each other at a type where the environment is allowed to operate on the queue as well. We also note that the converse of the axiom is obvious by stuttering and mumbling. After consuming a dummy argument $()$ let the resulting pilot trace be $(h_1, k_1) \dots (h_i, k_i) \dots (h_n, k_n)a$ and h'_1 be the start heap to match. We can now assume that the passages from k_i to h_{i+1} are according to the protocol, i.e. $k_i \xrightarrow{\text{msq}(X)} h_{i+1}$. Namely, should this not be the case we are free to make arbitrary moves and still win the game by default of the environment player. Therefore, there must exist i such that in the move (h_i, k_i) the element a is dequeued and $h_j = k_j$ holds for $j \neq i$. We can thus match this trace by a trace in the semantics of `atomic(dequeue ())` by stuttering until i :

$(h'_1, h'_1) \dots (h'_i, \dots$

where h_j and h'_j have the same content, but not necessarily the exact same layout. Given the environment’s allowed effects it is then clear that also h_i and h'_i have the same content, but not necessarily the same as h_1 and h'_1 because in the meantime other operations on the queue might have succeeded. We then dequeue the corresponding element from h'_i leading to k'_i and continue by stuttering.

$\dots, k'_i)(h'_{i+1}, h'_{i+1}) \dots (h'_n, h'_n)a'$

It is now clear that this is a matching trace and that $a = a'$ so we are done.

Notice that the congruence rules now allow us to deduce the equivalence of $op_1 || \dots || op_n$ and `atomic(op_1)` $|| \dots ||$ `atomic(op_n)` for op_i being enqueues or dequeues, which effectively amounts to linearizability.

10. Discussion

We have shown how a simple effect system for stateful computation and its relational semantics, combined with the notion of abstract locations, scales to a concurrent setting. The resulting type system provides a natural and useful degree of control over the otherwise anarchic possibilities for interference in shared variable languages, as demonstrated by the fact that we can delineate and prove the conditions for non-trivial contextual equivalences, including fine-grained data structures.

The primary goal of this line of work is not so much to find reasoning principles that support the most subtle equivalence arguments for particular programs, but rather to capture more generic properties of modules, expressed in terms of abstract locations and relatively simple effect annotations, that can be exploited by clients (including optimizing compilers) in external reasoning and transformations. But there are of course, particularly in view of the fact that we allow deeper reasoning to be used to establish that expressions can be assigned particular effect-refined types, very close connections with other work on richer program logics and models.

Rely-guarantee reasoning is widely used in program logics for concurrency, including relational ones [18], whilst our abstract locations are very like the *islands* of Ahmed et al [4]. Recent work of Turon et al [27] on relational models for fine-grained concurrency introduces richer abstractions, notably state transition systems expressing inter-thread protocols that can involve ownership transfer. These certainly allow the verification of more complex fine-grained algorithms than can be dealt with in our setting, and it would be natural to try defining an effect semantics over such a model. Indeed, one might reasonably hope that effects could provide something of a ‘simplifying lens’, with refined types capturing things that would otherwise be extra model structure or more complex invariants, such that the combination does not lead to further complexity. The use of Brookes’s trace model (also used by, for example, Turon and Wand [28]) already seems to bring some simplification compared to transition systems or resumptions.

Birkedal et al [10] have also studied relational semantics for effects in a concurrent language. The language considered there has dynamic allocation via regions and higher-order store, neither of which we have here. On the other hand, their invariants are based on simply-typed concrete locations and thus do not allow to capture effects at the level of whole datastructures as abstract locations do. As a result, the examples in [10] are of a simpler nature than ours. Furthermore, we offer a subtler parallelization rule, distinguish transient and end-to-end effects, and validate other effect-dependent equivalences like commuting, lambda hoist, deadcode and duplication. Our use of denotational methods and in particular the extension of Brookes’ trace semantics to higher-order functions does result in a rather simpler and more intuitive definition of the logical relation by comparison with [10]. While some of the complications are due to the dynamic allocation and typed locations, others like the explicit step counting, the need for effect-instrumented operational semantics, and the separation of branches in the definition of safety are not. We thus see our work also as a proof-of-concept for denotational semantics in the realm of higher-order concurrent programming.

The ‘RGSim’ relation proposed by Liang *et al.* for proving concurrent refinements under contextual assumptions also has many similarities with our logical relation [18, Def.4]. The focus of that work is on proving particular equivalences and refinements, whereas we encapsulate general patterns of behaviour in a refined

type system and can show the soundness of generic program transformations relying only on effect types (which combine smoothly with hand proofs of particular equivalences).

There are many directions for further work. Most importantly, we would like to add dynamic allocation of abstract locations following [6]. In addition to relieving us from having to set up all data structures in the initial heap this would, as we believe, also allow us to model and reason about lock-based protocols in an elegant way. Other possible extension include higher-order store and weak concurrency models.

References

- [1] M. Abadi, C. Flanagan, and S. N. Freund. Types for safe locking: Static race detection for java. *ACM Trans. Program. Lang. Syst.*, 28(2):207–255, 2006.
- [2] M. Abadi and L. Lamport. The existence of refinement mappings. *Theor. Comput. Sci.*, 82(2):253–284, 1991.
- [3] S. Abramsky and A. Jung. Domain theory, 1994. Online Lecture Notes, available from CiteSeerX.
- [4] A. Ahmed, D. Dreyer, and A. Rossberg. State-dependent representation independence. In *POPL*, 2009.
- [5] T. Amtoft, F. Nielson, and H. R. Nielson. *Type and Effect Systems: Behaviours for Concurrency*. World Scientific, 1999.
- [6] N. Benton, M. Hofmann, and V. Nigam. Abstract effects and proof-relevant logical relations. In *POPL*, pages 619–632, 2014.
- [7] N. Benton, A. Kennedy, L. Beringer, and M. Hofmann. Relational semantics for effect-based program transformations with dynamic allocation. In *PPDP*, 2007.
- [8] N. Benton, A. Kennedy, M. Hofmann, and L. Beringer. Reading, writing and relations: Towards extensional semantics for effect analyses. In *APLAS*, volume 4279 of *LNCS*, 2006.
- [9] N. Benton, A. Kennedy, and G. Russell. Compiling Standard ML to Java bytecodes. In *ICFP*, 1998.
- [10] L. Birkedal, F. Sieczkowski, and J. Thamsborg. A concurrent logical relation. In P. Cégielski and A. Durand, editors, *CSL*, volume 16 of *LIPICs*, pages 107–121. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2012.
- [11] L. Birkedal, M. Tofte, and M. Vejstrup. From region inference to von Neumann machines via region representation inference. In *Proceedings of the 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '96)*, 1996.
- [12] N. Broberg and D. Sands. Flow locks: Towards a core calculus for dynamic flow policies. In *15th European Symposium on Programming (ESOP '06)*, volume 3924 of *LNCS*. Springer, 2006.
- [13] S. D. Brookes. Full abstraction for a shared-variable parallel language. *Inf. Comput.*, 127(2):145–163, 1996.
- [14] J. W. Coleman and C. B. Jones. A structural proof of the soundness of rely/guarantee rules. *J. Log. Comput.*, 17(4):807–841, 2007.
- [15] C. Flanagan and S. Qadeer. A type and effect system for atomicity. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '03)*, 2003.
- [16] D. K. Gifford and J. M. Lucassen. Integrating functional and imperative programming. In *LISP and Functional Programming*, 1986.
- [17] O. Kammar and G. D. Plotkin. Algebraic foundations for effect-dependent optimisations. In *POPL*, 2012.
- [18] H. Liang, X. Feng, and M. Fu. A rely-guarantee-based simulation for verifying concurrent program transformations. In J. Field and M. Hicks, editors, *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012*, pages 455–468. ACM, 2012.
- [19] N. A. Lynch and F. W. Vaandrager. Forward and backward simulations, ii: Timing-based systems. *Inf. Comput.*, pages 1–25, 1996.
- [20] M. M. Michael and M. L. Scott. Nonblocking algorithms and preemption-safe locking on multiprogrammed shared memory multiprocessors. *J. Parallel Distrib. Comput.*, 51(1):1–26, May 1998.
- [21] N. Benton and P. Buchlovsky. Semantics of an effect analysis for exceptions. In *3rd ACM Workshop on Types in Language Design and Implementation (TLDI '07)*, 2007.
- [22] R. D. Nicola and M. Hennessy. Testing equivalence for processes. In *ICALP*, pages 548–560, 1983.
- [23] F. Pessaux and X. Leroy. Type-based analysis of uncaught exceptions. In *Proceedings of the 26 ACM Symposium on Principles of Programming Languages (POPL '99)*, 1999.
- [24] G. D. Plotkin. A powerdomain construction. *SIAM J. Comput.*, 5(3):452–487, 1976.

- [25] J. Thamsborg and L. Birkedal. A Kripke logical relation for effect-based program transformations. In *ICFP*, 2011.
- [26] J.-B. Tristan and X. Leroy. A simple, verified validator for software pipelining. In *POPL*, 2010.
- [27] A. J. Turon, J. Thamsborg, A. Ahmed, L. Birkedal, and D. Dreyer. Logical relations for fine-grained concurrency. In R. Giacobazzi and R. Cousot, editors, *POPL*, pages 343–356. ACM, 2013.
- [28] A. J. Turon and M. Wand. A separation logic for refining concurrent objects. In T. Ball and M. Sagiv, editors, *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, pages 247–258. ACM, 2011.

A. Proof of Theorem 7.7

Proof. In each case, using Corollary 3.2 and Lemma 3.3 (for case 6), we can in fact assume w.l.o.g. that the assumed pairs are in $T_0(\dots)$ rather than $T(\dots)$.

Ad 1. Let $(t, a) \in q_i(U)$, i.e. $a = p_i(a_0)$ where $(t, a_0) \in U$. By down-closure ([Down]) we also have $(t, a) \in U$. We can now play the strategy guaranteed by the assumption $(U, U') \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ which will yield (depending on the opponent's moves) a trace t' and a value a' such that $(t', a') \in U'$ and $(p_i(a), a') \in E$. Now, since E is a specification we get $(p_i(a), p_i(a')) \in E$ noting that p_i is idempotent. So, we modify the strategy so as to return $p_i(a')$ rather than a' and thus obtain a winning strategy asserting the desired conclusion.

Ad 2 This is an easy consequence from 1.

Ad 3 Pick $(U, U') \in T_0(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$. Since $T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ is closed under suprema it suffices to show that $(q_j(U^\dagger), q_j(U'^\dagger)) \in T(E, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ for each j . Fix such j and pick $(t, p_j(a)) \in q_j(U^\dagger)$, thus $(t, a) \in U^\dagger$.

By induction on the closure process we can assume w.l.o.g. that (t, a) arises from $(t_1, a) \in U$ by a single mumbling or stuttering step or that $(t, a_1) \in U$ for some $a_1 \geq a$ or else that $(t, a_i) \in U$ where $\sup_i a_i = a$.

In the former two cases fix a strategy for the original element of U . We will use this strategy to build a new one demonstrating that $(t, a) \in U'$, hence $(t, p_j(a)) \in q_j(U')$ as required.

If (t, a) arises by stuttering, so $t = u(h, h)v$ and $t_1 = uv$ we play the strategy until u is worked off. If the opponent then produces a heap h' to match h we answer h' .

Now $[\varepsilon_1](h, h', h, h')$ is always true (Lemma 6.2) so this is a legal move. Thereafter, we continue just as in the original strategy. In the special case where v is empty, we must also show that $[\varepsilon_3](h_1, h'_1, h, h')$ knowing $[\varepsilon_3](h_1, h'_1, k_n, k'_n)$ where $u = (h_1, k_1) \dots (h_n, k_n)$ and $u' = (h'_1, k'_1) \dots (h'_n, k'_n)$ is the matching trace. We have $[\varepsilon_2](k_n, k'_n, h, h')$ for otherwise opponent's playing h' would have been illegal. Since, by assumption $\varepsilon_2 \subseteq \varepsilon_3$, we can conclude $[\varepsilon_3](k_n, k'_n, h, h')$ and then $[\varepsilon_3](h_1, h'_1, h, h')$ by Lemma 6.2(3&1).

If (t, a) arises by mumbling then we must have $t = u(h_1, h_3)v$ and $t_1 = u(h_1, h_2)(h_2, h_3)v$. We play until the strategy has produced a match h'_2 for h_2 . So far, the play has produced a trace u' matching u , and a state h'_1 so that $[\varepsilon_1](h_1, h'_1, h_2, h'_2)$. Now, we can ask what the original strategy would produce if we gave it (temporarily assuming opponent's role) the state h'_2 as a match for h_2 . Note that this is legal because $[\varepsilon_2](h_2, h'_2, h_2, h'_2)$. The strategy will then produce h'_3 such that $[\varepsilon_1](h_2, h'_2, h_3, h'_3)$ and our answer in the play on the new trace against the challenge h'_1 will be this very h'_3 . Indeed, by composing tiles (Lemma 6.2) we have $[\varepsilon_1](h_1, h'_1, h_3, h'_3)$ as required. Thereafter, the play continues according to the original strategy.

For down-closure, we play the strategy against (t, a_1) yielding a match $(t', a'_1) \in U'$ where $a_1 E a'_1$. That same strategy also wins against (t, a) because $a E a'_1$ since E is a value specification.

For closure under $[\text{Sup}]$, finally, pick i so that $a_i \geq p_j(a)$ recalling that $a = \sup_i a_i$. Since we have a winning strategy for (t, a_i) , we also have one (by down-closure which was already proved) for $(t, p_j(a))$ as required.

Ad 4. Suppose $a E a'$. By 3 which we have just proved we only need to match elements of the form $((h, h)a)$. The opponent plays h' where $h \stackrel{\text{rds}(\varepsilon_3)}{\sim} h'$ and we answer with h' itself and a' . This is always a legal move (Lemma 6.2) and $a E a'$, so we win the game.

Ad 5. Again, we only need to match traces of the form $((h, h_1), a)$ where $c(h) = (h_1, a)$. In this case, suppose that the opponent

plays h' where $h \stackrel{\varepsilon_3}{\sim} h'$. The assumption gives (h'_1, a') such that $c'(h') = (h'_1, a')$ and $[\varepsilon_1](h, h', h_1, h'_1)$ and aEa' . We thus play h'_1 and a' and indeed $[\varepsilon_{1/3}](h, h', h_1, h'_1)$ and aEa' hold so this is a winning move.

Ad 6. Suppose $(f, f') \in E_1 \rightarrow T(E_2, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ and $(U, U') \in T(E_1, \varepsilon_1, \varepsilon_2, \varepsilon_3)$. Suppose that $(uv, b) \in ap(f, U)$ where $(u, a) \in U$ and (v, b) in $f(a)$ (note that we can ignore the \dagger -closure). We need to produce a trace $(u'v', b') \in ap(f', U')$ such that $(u', a') \in U'$ and (v', b') in $f'(a')$ and bE_2b' . Assume that:

$$u = (h_1, k_1) \cdots (h_n, k_n) \text{ and } v = (h_{n+1}, k_{n+1}) \cdots (h_{n+m}, k_{n+m})$$

We are given a heap h'_1 , such that $h_1 \stackrel{\text{rds}(\varepsilon_3)}{\sim} h'_1$. We can use the strategy S_1 from $(U, U') \in T(E_1, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ for (u, a) . We play according to S_1 to work off the u -part. This results in a matching trace $u' \in U'$:

$$u' = (h'_1, k'_1) \cdots (h'_n, k'_n)$$

where $[\varepsilon_3](h_1, h'_1, k_n, k'_n)$ and $(a, a') \in E_2$. We get $(f(a), f(a')) \in T(E_2, \varepsilon_1, \varepsilon_2, \varepsilon_3)$. Now, we are given a heap h'_{n+1} that is an environment move forming the tile $[\varepsilon_2](k_n, k'_n, h_{n+1}, h'_{n+1})$. From the fact that $\varepsilon_2 \subseteq \varepsilon_3$ and Lemma 6.2(5) we can conclude $h_{n+1} \stackrel{\text{rds}(\varepsilon_3)}{\sim} h'_{n+1}$.

Thus we can continue our play by using the strategy S_2 from $(f(a), f(a')) \in T(E_2, \varepsilon_1, \varepsilon_2, \varepsilon_3)$ which yields a continuation v' of our trace and a final answer b' . It is then clear that $(u'v', b') \in \text{bnd}(f', U')$ so this combination of strategies does indeed win.

Ad 7. Suppose that $(U_1, U'_1) \in T(E_1, \varepsilon_1, \varepsilon \cup \varepsilon_2, \varepsilon \cup \varepsilon_2 \cup \varepsilon')$ and $(U_2, U'_2) \in T(E_2, \varepsilon_2, \varepsilon \cup \varepsilon_1, \varepsilon \cup \varepsilon_1 \cup \varepsilon')$ and let $(t, (a, b)) \in U_1 \mid U_2$, thus $\text{inter}(t_1, t_2, t)$ (ignoring \dagger by item 3) where $(t_1, a) \in U_1$ and $(t_2, b) \in U_2$. Let S_1, S_2 be corresponding winning strategies. The idea is to use S_1 when we are in t_1 and to use S_2 when we are in t_2 . Supposing that t starts with a t_1 fragment we begin by playing according to S_1 . Let t be of the form:

$$t = (h_1, k_1) \cdots (h_n, k_n)(h_{n+1}, k_{n+1}) \cdots (h_{n+m}, k_{n+m}) \\ (h_{n+m+1}, k_{n+m+1}) \cdots (h_{n+m+k}, k_{n+m+k}) \cdots (h_p, k_p)$$

composed of pieces of the traces t_1 and t_2 . Assume w.l.o.g. that the first piece $(h_1, k_1) \cdots (h_n, k_n)$ is a part of t_1 . We are given an initial heap h'_1 such that $h \stackrel{\text{rds}(\varepsilon \cup \varepsilon' \cup (\varepsilon_1 \cup \varepsilon_2))}{\sim} h'$. Since $\text{rds}(\varepsilon_1 \cup \varepsilon_2) = \text{rds}(\varepsilon_1) \cup \text{rds}(\varepsilon_2)$, we can apply strategy S_1 to guide us through the first part of the game, obtaining:

$$(h'_1, k'_1) \cdots (h'_n, k'_n)$$

Moreover, we have an environment move which forms the tile $[\varepsilon](k_n, k'_n, h_{n+1}, h'_{n+1})$. Thus, we have the tile $[\varepsilon \cup \varepsilon_1](h_1, h'_1, h_{n+1}, h'_{n+1})$ which can be seen as an environment move for t_2 . Therefore, we can use strategy S_2 for the U' and continue the game, obtaining the trace piece:

$$(h'_{n+1}, k'_{n+1}) \cdots (h'_{n+m}, k'_{n+m})$$

Now, we can return to the S_1 game as the trace above is seen as an environment move for U . Alternating these strategies, we get a trace t which is in $(U \mid U')$. Let (a', b') be the final values reached at the end. It is clear that $[\varepsilon \cup \varepsilon' \cup \varepsilon_1 \cup \varepsilon_2](h, h', h_p, h'_p)$ and also aE_1a' and bE_2b' .

It remains to assert the stronger statement $[\varepsilon \cup \varepsilon' \cup (\varepsilon_1 \cup \varepsilon_2)](h, h', h_p, h'_p)$. To see this suppose that $wr_1 \in \varepsilon_1 \setminus \varepsilon_2 \setminus \varepsilon \setminus \varepsilon'$. Since the entire game can be viewed as an instance of the game U_1 vs U'_1 with interventions by U_2 vs. U'_2 regarded as environment interactions we have $[\varepsilon \cup \varepsilon_2 \cup \varepsilon'](h, h', h_p, h'_p)$ so that in fact $h \stackrel{1}{=} h_p$ and $h' \stackrel{1}{=} h'_p$. The case of co_1 and $\varepsilon_1, \varepsilon_2$ interchanged is analogous.

Ad 8. This is direct from the definition of atomic and appealing on the fact that $(U, U') \in T(E, \varepsilon_1, \emptyset, \varepsilon_3)$. \square

B. Proof of Theorem 9.1

Proof. Commuting. By Theorem 7.7(3) we can assume our pilot trace t to be of the form:

$$(h_1, k_1)(h_2, k_2) \cdots (h_n, k_n) (h_{n+1}, k_{n+1}) \cdots (h_{n+m}, k_{n+m}) (a, b)$$

where

$$t_1 = (h_1, k_1)(h_2, k_2) \cdots (h_n, k_n) \quad v_1 \in U_1 \\ t_2 = (h_{n+1}, k_{n+1}) \cdots (h_{n+m}, k_{n+m}) \quad v_2 \in U_2$$

We make similar use of Theorem 7.7(3) in the subsequent cases without explicit mention.

We are also given a heap h'_1 such that $h_1 \stackrel{\text{rds}(\varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2)}{\sim} h'_1$. Because $\varepsilon'_1 \perp \varepsilon'_2$, h_1 and h_{n+1} agree on the reads of ε'_2 . Thus we can start a game U_2 vs. U'_2 using h'_1 and t_2 . We forward all environment's moves from the main game to the side game and use the responses from the side game to answer in the main game. Suppose that the side game leads to the valid U_2 -trace

$$(h'_1, k'_1)(h'_2, k'_2) \cdots (h'_m, k'_m) v'_2$$

where $v_2 E_2 v'_2$ and (1) $[\varepsilon^C \cup \varepsilon'_2](h_{n+1}, h'_1, k_{n+m}, k'_m)$. Notice that in the global game these are legal responses as $[\varepsilon^C \cup \varepsilon'_2](h_i, h'_i, k_i, k'_i)$ for $1 \leq i \leq m$.

We now have an environment move $[\varepsilon](k_m, k'_m, h_{m+1}, h'_{m+1})$. Since $\varepsilon'_1 \perp \varepsilon$ and $\varepsilon'_2 \perp \varepsilon'_1$, the heaps h'_1 and h'_{m+1} agree in the reads of ε'_1 . Therefore, we can run a game U_1 vs. U'_1 using h'_{m+1} and t_1 , obtaining the trace:

$$(h'_{m+1}, k'_{m+1})(h'_{m+2}, k'_{m+2}) \cdots (h'_{m+n}, k'_{m+n}) v'_1$$

where $v_1 E_1 v'_1$ and (2) $[\varepsilon^C \cup \varepsilon'_1](h_1, h'_{m+1}, k_n, k'_{m+n})$. The reasoning is similar to the use of the previous game.

Thus we have that $(v_1, v_2)(E_1 \times E_2)(v'_1, v'_2)$.

Now, we need to conclude that $[\varepsilon^C \cup \varepsilon'_1 \cup \varepsilon'_2](h_1, h'_1, k_{n+m}, k'_{m+n})$. This follows from the fact that $\varepsilon'_1 \perp \varepsilon'_2$ and (1) and (2). In particular, from (1) and $\varepsilon'_1 \perp \varepsilon'_2$, we get that k_{m+n} and k'_{m+n} agree on the locations in ε'_2 , while from (2), we get that k_{m+n} and k'_{m+n} agree on the locations in ε'_1 . This finishes the proof.

Duplicated. Assume given a trace in U :

$$t = (h_1, k_1) \cdots (h_n, k_n) v$$

and a heap h'_1 such that $h_1 \stackrel{\text{rds}(\varepsilon_2 \cup \varepsilon')}{\sim} h'_1$. Since $\varepsilon_2 \perp \varepsilon_1$ and $\text{rds}(\varepsilon') \cap \text{wr}(\varepsilon') = \emptyset$, we have that h_1 and k_n agree on the reads of ε' .

We start by simply stuttering:

$$t' = (h'_1, h'_1)(h'_2, h'_2) \cdots (h'_n, ??)$$

where $[\varepsilon](k_i, h_{i+1}, k'_i, h'_{i+1})$ for $1 \leq i \leq n + m$. Notice that for $1 \leq i \leq n - 1$, we have $[\varepsilon^C_1](h_i, h'_i, h_{i+1}, h'_i)$. So the stuttering moves are valid responses.

We will now play U_1 vs. U'_1 to construct the missing heap “?”. We first run a game using h'_n and t , where the environment moves are simply stutter moves:

$$(h'_n, q_1)(q_1, q_2) \cdots (q_{n-1}, q_n) v'_1$$

such that $v E v'_1$ and $[\varepsilon^C \cup \varepsilon'](h_1, h'_n, k_n, q_n)$. Notice that using stuttering environment moves are valid as $[\varepsilon^C](k_i, q_i, h_{i+1}, q_i)$ for $1 \leq i \leq n - 1$.

Since h_1 and k_n agree on the reads of ε' and q_n and k_n agree on $\text{rds}(\varepsilon')$ from $[\varepsilon^C \cup \varepsilon'](h_1, h'_n, k_n, q_n)$, we can run the game U_1 vs. U'_1 again on q_n and t with stutter environment moves:

$$(q_n, q_{n+1})(q_{n+1}, q_{n+2}) \cdots (q_{n+m-1}, q_{n+m}) v'_2$$

where $v E v'_2$ and $[\varepsilon^C \cup \varepsilon'](h_1, q_n, k_n, q_{n+m})$. Thus, $(v, v')(E \times E)(v'_1, v'_2)$.

We now put $?? := q_{m+n}$ which leads to a valid trace due to repeated mumbling. Finally, $[\varepsilon \cup \varepsilon'_2](h_1, h'_1, k_n, q_{n+m})$ follows from $[\varepsilon^C \cup \varepsilon'](h_1, q_n, k_n, q_{n+m})$ and $\varepsilon \perp \varepsilon'$.

Pure. We start with a trace from $rn(v)$, for example $(h_1, h_1), v$ and an arbitrary heap h'_1 . We now consider the game involving U vs. U' on t, v and h'_1 :

$$\begin{aligned} t &= (q_1, k_1)(q_2, k_2) \cdots (q_n, k_n), v \\ t' &= (h'_1, k'_1)(k'_1, k'_2) \cdots (k'_{n-1}, k'_n), v' \end{aligned}$$

We have that vEv' and $[\varepsilon_3](q_1, h'_1, k_n, k'_n)$. By mumbling, $(h'_1, k'_n) \in U'$. We can reply with k'_n in the main game.

Dead. Assume given a trace of the form:

$$(h_1, k_1) \cdots (h_n, k_n) v$$

and h'_1 such that $h_1 \stackrel{\text{rds}(\varepsilon_3)}{\sim} h'_1$. We now initiate a side game U vs. U' on this trace and respond in the main game by stuttering. Thus, we obtain traces $(h'_1, h'_1) \cdots (h'_n, h'_n) ()$ in the main game and $(h'_1, k'_1) \cdots (h'_n, k'_n) v'$ in the side game.

The main trace is in $rn()$. The side game tells us that $v = ()$ and that $h_i \xrightarrow{\varepsilon_1} k_i$ and therefore $[\varepsilon_1^C](h_i, h'_i, k_i, h'_i)$. It remains to show that $[\varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2](h_1, h'_1, k_n, k'_n)$. This follows from the fact that ε_1 has only reads as h_i and k_i agree on all locations.

Parallelization. We start with a trace in $U_1 \parallel U_2$. Assume that the trace is of the following form:

$$t_{1,1}t_{2,1}t_{1,2}t_{2,2} \cdots t_{1,n}t_{2,n} (v_1, v_2)$$

where each $t_{i,j}$ is a possibly empty sequence of moves of the form $(h_{i,j}^1, k_{i,j}^1) \cdots (h_{i,j}^{m_{i,j}}, k_{i,j}^{m_{i,j}})$ and

$$\begin{aligned} t_1 &= t_{1,1} \cdots t_{1,n} \quad v_1 \in U_1 \\ t_2 &= t_{2,1} \cdots t_{2,n} \quad v_2 \in U_2 \end{aligned}$$

are traces from U_1 and U_2 , respectively. We are also given a heap h'_1 such that $h_{1,1}^1 \stackrel{\text{rds}(\varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2)}{\sim} h'_1$. We also have $h_{1,1}^1 \stackrel{\text{rds}(\varepsilon^C \cup \varepsilon_2^C \cup \varepsilon'_1)}{\sim} h'_1$. We run a side game U_1 vs. U'_1 using h'_1 and t_1 , yielding:

$$t'_{1,1} \cdots t'_{1,n} v'_1$$

Assume that (h'_1, k'_1) and (h'_o, k'_o) are, respectively, the first and last moves of this trace. We have $v_1 E_1 v'_1$ and (1) $[\varepsilon^C \cup \varepsilon_2^C \cup \varepsilon'_1](h_{1,1}^1, h'_1, k_{1,n}^m, k'_o)$. Notice that these are legal moves in the global game as we have $[\varepsilon_1^C \cup \varepsilon_2^C]$ tiles for the player moves and $[\varepsilon]$ times for the environment moves.

Now, assume there is an environment move (k_o, h'_{o+1}) . Since $\varepsilon_1 \perp \varepsilon_2$ and $\varepsilon \perp \varepsilon_2$, the heaps $h_{1,1}^1$ and $h_{2,1}^1$ agree on the reads of ε'_2 and h'_1 and h'_{o+1} also agree on the reads of ε'_2 . (Notice as well that $\text{wrs}(\varepsilon_1) \cap \text{rds}(\varepsilon'_2) = \emptyset$ as $\varepsilon^C \cup \varepsilon_1^C \cup \varepsilon_2$ is a valid effect.) Therefore, we can invoke an U_2 game using h'_{o+1} and t_2 , obtaining the trace:

$$t'_{2,1} \cdots t'_{2,n} v'_2$$

Assume that (h'_{o+1}, k'_{o+1}) and (h'_{o+p}, k'_{o+p}) are, respectively, the first and last moves of this trace. We have $v_2 E_2 v'_2$ and (2) $[\varepsilon^C \cup \varepsilon_1^C \cup \varepsilon'_2](h_{2,1}^1, h'_{o+1}, k_{2,n}^m, k'_{o+p})$. For the same reasons as above, these are legal moves in the global game.

Therefore $(v_1, v_2)(E_1 \times E_2)(v'_1, v'_2)$.

We need now to prove that $[\varepsilon \cup \varepsilon'_1 \cup \varepsilon'_2](h_{1,1}^1, h'_1, k_{2,n}^m, k_{o+p})$. From (1) and $\varepsilon_1 \perp \varepsilon_2$ and $\varepsilon \perp \varepsilon_1$, we have that $k_{2,n}^m$ and k_{o+p} agree on the locations of ε_1 . Similarly, $k_{2,n}^m$ and k_{o+p} agree on the locations of ε_2 . Since there are only ε tiles and $\varepsilon \perp \varepsilon_1$ and $\varepsilon \perp \varepsilon_2$, $k_{2,n}^m$ and k_{o+p} agree on the locations of ε . This finishes the proof. \square